



Visualizing Modular Forms

David Lowry-Duda


University of Oregon

17 May 2021

ICERM and Brown University


Plotting functions is easy to do, and hard to do well. There are many plotting programs and it is often pretty easy to produce a plot that represents *all* of the information about a function.

But many plots can misrepresent the underlying data and mislead the viewer. Different plots and different choices appeal to different aspects of the viewer's intuition — and it's important to recognize that these defaults are probably not always right.



The humble sine plot,
according to python's
matplotlib.

The sine plot in
sagemath.



I began investigating visualizations of modular forms in early 2019. The LMFDB¹ was adding a large number of (classical holomorphic cuspidal) modular forms and their L -functions to the database. We thought: could we *show* the modular form too?

I'd never seen a plot of a modular form. Did I know what a modular form looked like?

At the fall 2019 ICERM program on Illustrating Mathematics, I talked to many others about this problem. I should thank Frank Farris and Elias Wegert for introducing me to the world of complex function visualization.

¹L-function and modular form database, at LMFDB.org

Plotting complex-valued functions is hard. The graph of a function

$$f : \mathbb{C} \longrightarrow \mathbb{C}$$

is fundamentally 4-dimensional. Representing 3-dimensional objects on a 2-dimensional medium is hard enough. 4 is outright challenging.

It is necessary to either throw away information to reduce the dimension, or to rely on non-spacial ways of representing data, or to use a mixture of both.

This is the default plot of the identity function in sage. Given a point $z = re^{i\theta}$, we let the color be determined by the argument θ and the brightness be determined by the magnitude r .

The argument is periodic and the color wheel is circular, so this representation has some logic to it.

Brightness for magnitude isn't very effective for strict comparison.



Under the hood, many decisions have already been made.

How do we map from angles to colors? Are these the right colors to use?

How do we map from magnitude to brightness?

For magnitude to brightness, the defaults need to work for a generic function, even if the function grows enormously. In sage, this is done with the map

$$\mathbb{C} \longrightarrow \text{Brightness}$$

$$z \mapsto \frac{4}{\pi} \arctan(\log(\sqrt{|z|} + 1)) - 1.$$

This gives a number between -1 (black) and 1 (white). Values near 0 are darker, and values near ∞ are lighter.



If we reverse this, so that numbers nearer to 0 are lighter, we get the first plot on the left.

If we also change the magnitude-to-brightness map (here, I removed the square root), we get the second plot on the left.

None of these are obviously better than the others. But these are just a few examples of the knobs we can twist when examining plots of complex functions.

Modular Forms

A modular form is a well-behaved complex-valued function f defined on the upper half-plane \mathcal{H} . These forms satisfy symmetries of the form

$$f\left(\frac{az+b}{cz+d}\right) = (cz+d)^k f(z), \quad \gamma = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \Gamma(N) < \mathrm{SL}(2, \mathbb{Z})$$

for an integer k (the *weight*) and an integer N (the *level*). Here, $\gamma z = \frac{az+b}{cz+d}$ is a group action of $\mathrm{SL}(2, \mathbb{Z})$ on \mathcal{H} .

This implies that these forms will be periodic horizontally,

$$f(z+N) = f(z),$$

and satisfy a more complicated periodic condition in other directions.

Our fundamental examples of modular forms are holomorphic and cuspidal. Holomorphic cuspforms vanish as $\text{Im } z \rightarrow \infty$ or as z approaches any rational point on the real line. Each such form has a Fourier expansion of the shape

$$f(z) = \sum_{n \geq 1} a(n) e^{2\pi i n z / h}$$

for some $h \mid N$.

In the Fourier expansion, the weight manifests as a bound on the coefficients

$$a(n) \ll n^{\frac{k-1}{2} + \epsilon},$$

and the level manifests in the normalizing factor h . Of course, both the weight and level affect the symmetries satisfied by the expansion (but this is fairly opaque).

After we compute the Fourier coefficients $a(n)$ (or more typically, look them up in the LMFDB or use CAS like sage/pari-gp/magma to compute them for us), we can explicitly compute the form.

So how do we visualize the form?

We'll begin with the unique holomorphic cusp form of weight 12 on $SL(2, \mathbb{Z})$ (so level 1), the Δ function famously studied by Ramanujan (and elliptic function specialists in the 1800s). This is given by

$$\begin{aligned}\Delta(z) &= q \prod_{n \geq 1} (1 - q^n)^{24}, \quad (q = e^{2\pi iz}) \\ &= q - 24q^2 + 252q^3 - 1472q^4 + 4830q^5 + \cdots\end{aligned}$$

In many ways, this is the easiest modular form to study.



Plotting $\Delta(z)$ on the whole upper-plane is not an option. But it might be reasonable to plot one horizontal period. Further, since $\Delta(z) \rightarrow 0$ as $\text{Im } z \rightarrow \infty$, we might hope that truncating our plot at some height Y omits little interesting behavior.

Plotting Δ in the on a small domain of \mathcal{H} looks like the following.

Δ plotted on $[-.5, 1.5] \times [0, 1.5]$
in the same colors as sage
defaults.

The horizontal periodicity is
clear. Other symmetries are
clearer if we compare with a
tiling of the fundamental
domain.

Each (hyperbolic) triangle at
bottom right is a fundamental
domain for $SL(2, \mathbb{Z}) \backslash \mathcal{H}$.
Matrices relate the values of
 $\Delta(z)$ on each triangle to the
others.



Alternately, we could choose a different representation of \mathcal{H} . The upper half-plane is conformal to the Poincaré disk, so it is possible to map \mathcal{H} to the disk \mathbb{D} . I choose an orientation by choosing to map to the disk with


$$\mathcal{H} \longrightarrow \mathbb{D}$$

$$0 \mapsto -i,$$

$$i \mapsto 0,$$

$$\infty \mapsto i.$$

The effect here is to present a little *medallion* representing a modular form, and the apparent vertical orientation of the imaginary line in the medallion is the same as the orientation in \mathcal{H} .



This is a very visually appealing representation. It contains “more” information than the part of \mathcal{H} , but it is challenging to reason about the behavior.

But on the other hand, it's pretty dark. It turns out that vanishing as $\text{Im } z \rightarrow \infty$ (and to other cusps) leads to lots of dark spots.

I like the basic format of these visualizations, but the choices of color leave much to be desired. In these visualizations, the argument of the form is obvious, but the magnitude, though theoretically shown everywhere, is in practice not very apparent away from locations where the function is very small.

One idea is to separate these two concerns. We're visualizing the modular forms by choosing a region of the domain and coloring each point of the domain in some predictable way. The domain choices are clear: regions of the upper half plane for detail views, or the disk for a (hyperbolic) bird's eye view.

But we can choose different colors.

This method of representing a complex function $f(z)$ by coloring each point in the domain of f is called *domain coloring*. In general, given any color function $g : \mathbb{C} \rightarrow \text{Color}$, we can produce a domain coloring of a function f on some region Ω by coloring each point $z \in \Omega$ by the color of $g(f(z))$.

Stated differently: we first choose a coloring of the complex plane, and then we pull back this coloring by the complex function f .

To understand what a domain coloring means, it's useful to see the color function.

All is much clearer through examples. Let us empirically study both color functions and a variety of modular forms.

We'll examine the following four forms repeatedly:


1. $\Delta(z)$, the Ramanujan Delta function,
2. A modular form for today's date, $f(z) = f_{5.17}(z)$ of level 5 and weight 17,
3. The non-holomorphic, half-integral weight modular form $g(z)$, the unique weight $9/2$ cuspform on $\Gamma_0(4)$, and
4. A non-holomorphic Maass form $\mu(z)$ on $\mathrm{SL}(2, \mathbb{Z})$. This is a modular form that's an eigenfunction of the hyperbolic Laplacian with eigenvalue $\frac{1}{4} + (20.1066\dots)^2$.

We'll also examine the color function applied to the simple polynomial $z^2(z - 1 - i)$.



Default sagemath plots.





The color function.





$z^2(z - 1 - i)$.





$\Delta(z)$



$f_{5.17}$



$g(z)$




$\mu(z)$


Default sagemath plots, except with light meaning close to zero.





The color function.




$z^2(z - 1 - i)$.




$\Delta(z)$




$f_{5.17}$



$g(z)$



$\mu(z)$



Nonclassical modular forms: half-integral weight

We're repeatedly studying four different forms. Two are classical holomorphic cuspidal Hecke eigenforms — the nicest, best behaved forms. The other two are less well-behaved in ways that make them non-classical.

The form g is a weight $9/2$ modular form on $\Gamma_0(4)$. The exact definition of weight $9/2$ is technically annoying to write down, but g is a holomorphic function and g^2 is a weight 9 form. If we define

$$\theta(z) = \sum_{n \in \mathbb{Z}} e^{2\pi i n^2 z} = 1 + 2q + 2q^4 + \cdots,$$

then g transforms in the same way as $\theta(z)^9$.

To each full-integral weight cuspform f and half-integral weight cuspform g ,

$$f(z) = \sum_{n \geq 1} a(n)q^n, \quad g(z) = \sum_{n \geq 1} b(n)q^n,$$

we can associate a Dirichlet series

$$L(s, f) = \sum_{n \geq 1} \frac{a(n)}{n^s}, \quad L(s, g) = \sum_{n \geq 1} \frac{b(n)}{n^s}.$$

The modularity and cuspidality of f and g ensure that these both continue to entire complex functions on \mathbb{C} . If f is a Hecke eigenform, then $L(s, f)$ has an Euler product and is expected to satisfy a Riemann Hypothesis. But g doesn't have an Euler product and doesn't satisfy a Riemann Hypothesis.

At the level of graphs, though, there is little to distinguish plots of full integral weight and half-integral weight forms. The analytic differences between having RH and not having RH are too subtle to be evident from visual behavior.

Nonclassical modular forms: Maass forms

A Maass form (of weight 0) is a *real-analytic* automorphic eigenfunction of the hyperbolic Laplacian

$$\Delta = -y^2 \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right).$$

We denote the eigenvalue $\Delta z = \lambda z$, $\lambda = \frac{1}{4} + R^2$. Maass forms with trivial central character can be normalized to be real-valued, and this is true of μ .

Each Maass form has an expansion

$$\mu(z) = \sum_{n \in \mathbb{Z}, n \neq 0} a(n) \sqrt{y} K_{iR}(2\pi ny) e^{2\pi i n x},$$

where $K_\nu(x)$ is the K -Bessel function.

Maass forms appear as a natural object in the spectral decomposition of modular forms: they span the discrete spectrum of Λ . Despite being so fundamental, they're very challenging to compute and remain largely mysterious.

The lines where $\mu(z) = 0$ is called the *nodal line*. As we've already seen, this line splits the domain into *nodal regions*. We expect that the number of nodal regions in the fundamental domain tends to ∞ as the eigenvalue λ tends to infinity. Ghosh, Reznikov, and Sarnak have a conditional proof (on LH/RH), but this also remains mysterious!

(I chose a Maass form with relatively large eigenvalue so that it would have lots of nodal domains).

Throw away the argument


These pictures were interesting, but largely black or white. What do we care more about — the magnitude or the argument? (The answer is not obvious).

We can simplify the pictures by throwing away the argument entirely and plotting only the magnitude.



A colorfunction that depends only on magnitude, completely ignoring argument.





The color function.





$z^2(z - 1 - i)$.





$\Delta(z)$



$f_{5.17}$




$g(z)$



$\mu(z)$

Relying on brightness/grayscale leads to haunting plots that are perhaps somewhat inexpressive. We could also cycle through the color wheel.



A colorfunction that depends only on magnitude, completely ignoring argument. Magnitude is shown by cycling through colors. Consecutive red bands differ by 1 in magnitude.





The color function.





$z^2(z - 1 - i)$.





$\Delta(z)$



$f_{5.17}$



$g(z)$




$\mu(z)$


A closely related idea would be to change how colors relate to the magnitude. Instead of consecutive bands of the same color corresponding to values that differ by 1 (or any constant), we could choose to allow consecutive bands of the same color to correspond to values that double or halve. That is, points in one red band are a multiplicative factor different than points in adjacent red bands.

Whether it's a multiplicative factor larger or smaller is left ambiguous. It's necessary to determine this from context.



A colorfunction that depends only on magnitude, completely ignoring argument. Magnitude is shown by cycling through colors. Consecutive red bands differ by a multiplicative factor of 10.





The color function.





$z^2(z - 1 - i)$.





$\Delta(z)$



$f_{5.17}$



$g(z)$




$\mu(z)$

Keep only some magnitude information


Elias Wegert suggested that I do a different sort of plot. In the following plots, we keep the argument (plotted as color again) and change how we plot the magnitude. To avoid dark or bright spots, we plot contours.

Actually, we don't compute contours — everything that follows is computed by domain coloring. Pulling back the domain coloring by the modular form gives the effect of contours.



A colorfunction with “contours”, indicating magnitude like terrain. I want to emphasize that one doesn’t compute contours — instead, once you’ve chosen a color function, the contours come for free. These colors come from matplotlib’s *viridis* colormap.





The color function.





$z^2(z - 1 - i)$.





$\Delta(z)$



$f_{5.17}$




$g(z)$




$\mu(z)$





Even minimalist color functions can give interesting visualizations.



The color function.







$z^2(z - 1 - i)$.




$\Delta(z)$

$f_{5.17}$




$g(z)$





$\mu(z)$



The color function.





$$z^2(z - 1 - i).$$





$\Delta(z)$

$f_{5.17}$




$g(z)$



$\mu(z)$


(This is not the most useful color function, but it's amusing).





The color function (and then tiled repeatedly). This is my wife and I in Iceland. There are puffins behind us!



$z^2(z - 1 - i)$.
Stunningly unflattering.





$\Delta(z)$



$f_{5.17}$



$g(z)$



$\mu(z)$

With few exceptions, I've changed the overall style of the color function rather than the individual color choices themselves. For example, the default colormap uses a cyclic colormap that goes through red, yellow, green, blue, purple, and back to red. This is a circle in colorspace.

The default choice is a quirk of history. But for various reasons, it's actually a lousy choice of color. When I make plots of modular forms for others, I'll use colors based on *viridis* or *cividis*, `matplotlib` colormaps. For more, check out


1. <http://jakevdp.github.io/blog/2014/10/16/how-bad-is-your-colormap/>
2. (SciPy 2015: A Better Default Colormap For Matplotlib)
<https://www.youtube.com/watch?v=xAoljeRJ3IU>

We've seen many different visualizations. I hope it's clear that there is no clear *best* visualization — instead, different visualizations highlight different aspects.

Some are beautiful but relatively light on information. Some theoretically perfectly describe the form, but in practice are muddled.

If you come across good visualizations, let me know!

To end, I'll note that I have a project in the works involving a different type of visualization:



(An in-motion perspective of $\Delta(z)$ is at
<https://www.youtube.com/watch?v=s6sdEbGNdic>)

Thank you very much.

**Please note that these slides are available on
my website (davidlowryduda.com).**