Adaptive interference inference

View On GitHub

Download PDF

This project is maintained by AndrewGYork in the York lab, and was funded by Calico Life Sciences LLC

Micro-publication

What is the signal-to-noise ratio of phase-contrast imaging? Can it be infinite?

Andrew G. York[1*], Sanjay R. Varma[1]

[1]Calico Life Sciences LLC, South San Francisco, CA 94080, USA

[*]Permanent email: andrew.g.york+inference@gmail.com

[*]Institutional email: agy+inference@calicolabs.com

Abstract

I'd like to invite you to consider one of my favorite physics questions, about how much information we can squeeze out of optical measurements. Even if you don't know any physics, the question translates cleanly to a pure mathematical puzzle, or alternatively, a coding challenge that might be well suited to machine-learning approaches. This is an old question, but an important one: optical measurements are an essential tool in every field of science and industry, and any improvement to our ability to extract information from light would have wide-ranging impact.

Intended audience

Anyone who likes math puzzles, coding challenges, and/or physics questions. Click the appropriate link below to jump to the relevant section:

Click here if you like math	Click here if you like coding	Click here if you like physics

Peer review status

Pre-print published October 15, 2018 (This article is not yet peer-reviewed)

Cite as: doi:10.5281/zenodo.1463273

Note that this is a limited PDF or print version; animated and interactive figures are disabled. For the full version of this article, please visit https://andrewgyork.github.io/adaptive_interference_inference

A math puzzle

Let's play a game. I know a complex number x. Your job is to infer x. You know that $|x| \leq 1$. *

You may choose any two complex numbers, a and b. I will calculate $y = |ax + b|^2$. Then I'll draw randomly from a Poisson distribution with expectation value y, and tell you the (stochastic) result r.

We repeat this process as many times as you'd like: on round i, you choose a_i and b_i, I'll calculate y_i and I'll tell you r_i. Your "cost" for round i is the quantity $|a_i|^2$, and you have a "budget" of A. If you go over budget (meaning, if you choose a value a_N such that $A < \sum_i^N |a_i|^2$), I won't answer, and the game is over.

What algorithm should you follow to infer x as accurately** as possible, for a given value of A? Ideally, express this algorithm as a short Python/Numpy script, ideally in a GitHub repository that I can link to here. Feel free to invite anyone who's interested to play.

Let the best algorithm win!

- Can you do any better than the algorithm $a_0 = \sqrt{A}$, $b_0 = 0$?
- Can you do any better than an algorithm where a and b are chosen in advance, independently of r_k ? For example: $a_k = \sqrt{A/N}$ and $b_k = \sqrt{A/N}e^{2\pi ik/N}$

Notes:

* For Bayesians: I'll choose the magnitude $|x|$ from a uniform distribution between 0 and 1, and I'll choose the phase $\angle x$ from a uniform distribution between 0 and 2π. Bonus points if your algorithm works robustly (or can be easily adapted) if x is drawn from a different distribution.

** Note I didn't define "accuracy"! I'll leave this intentionally flexible. I suspect the "ideal" algorithm will depend on the choice of definition, but I'm indifferent as long as your definition of accuracy seems sane to me.

A coding challenge

I've implemented an 'oracle' that will play a game with you. You can find a copy of the Python 3 code in our `adaptive_interference_inference` GitHub repository: ./code/game.py

```
import numpy as np

class Oracle:
    def __init__(self, budget):
        assert budget > 0
        self.budget = budget
        self.total_cost = 0
        self.round = 0
        self.a_history = []
        self.b_history = []
        self.r_history = []
        self.game_over = False

        # Generate some secret numbers.
        # Don't peek!
        self._magnitude = np.random.random(1)
        self._phase = 2*np.pi*np.random.random(1)
        self._x = self._magnitude * np.exp(1j*self._phase)
        return None

    def ask(self, a, b):
        if self.game_over:
            return None
        self.total_cost += np.abs(a)**2
        if self.total_cost > self.budget:
            self.game_over = True
            return None
        y = np.abs(a*self._x + b)**2
        r = np.random.poisson(y)
        self.round += 1
        self.a_history.append(a)
        self.b_history.append(b)
        self.r_history.append(r)
        return r
```

To play this game, import `game.py`, choose a value for `budget`, and create an instance of the `Oracle` object. Try to infer the value of `Oracle._magnitude` and/or `Oracle._phase` as accurately as you can, using only the return values of calls to the `Oracle.ask(a, b)` method. You are allowed to inspect the code of `game.py`, but after you instantiate an `Oracle` object, you are *not* allowed to inspect its `Oracle._x`, `Oracle._magnitude`, or `Oracle._phase` attributes; that's cheating!***

Here's an example of how you might use this oracle to play the game:

```
import numpy as np
import game

print("An example game")
oracle = game.Oracle(budget=1000)
print("Budget:", oracle.budget, '\n')
num_phases = 5
for phase_angle in np.arange(0, 2*np.pi, 2*np.pi/num_phases):
    print('Round', oracle.round)
    a = np.sqrt(oracle.budget / num_phases) * (1 - 1e-12)

    b = a * np.exp(1j*phase_angle)
    print(' a: %07s    (intensity)\n'%('%0.2f'%(np.abs(a)**2)),
          '   %07s*pi (phase)'%('%0.2f'%(np.angle(a)/np.pi)))
    print(' b: %07s    (intensity)\n'%('%0.2f'%(np.abs(b)**2)),
          '   %07s*pi (phase)'%('%0.2f'%(np.angle(b)/np.pi)))
    response = oracle.ask(a, b)
    print('oracle.ask(a, b):', response)
    print('Total cost: %0.2f / %0.2f\n'%(oracle.total_cost,
                                          oracle.budget))
print("Game over.")
print("From these responses, how well can you infer 'x'?")
print("\nIf you'd made smarter choices of 'a' and 'b' on each round,\n",
      "could you do a better job inferring 'x'?", sep='')
```

Note that this is just an example of how to ask the oracle questions. I didn't specify a method to infer `Oracle._x` based on the oracle's responses to `Oracle.ask(a, b)`. That's the game!

What's the best algorithm you can implement to infer the value of `Oracle._x` as accurately** as possible, for a given value of `Oracle.budget`? Ideally, express this algorithm as a short Python/Numpy script, ideally in a GitHub repository that I can link to here. Feel free to invite anyone who's interested to play.

Notes:

*** For machine-learners: feel free to inspect `Oracle._x`, `Oracle._magnitude`, and/or `Oracle._phase` while you're training and validating a model, just be careful that your *model* doesn't consult these values.

A physics question

Suppose we want to measure the complex optical transmission coefficient x of a thin uniform flat-faced slab of partially transparent unknown material[†]. When we transmit a collimated beam of monochromatic light through the slab, the slab reduces the beam's amplitude by $|x|$, and shifts the beam's phase by $\angle x$. How accurately can we measure this change in intensity and phase?

Consider the ideal interferometric measurement shown in Figure 1 below:

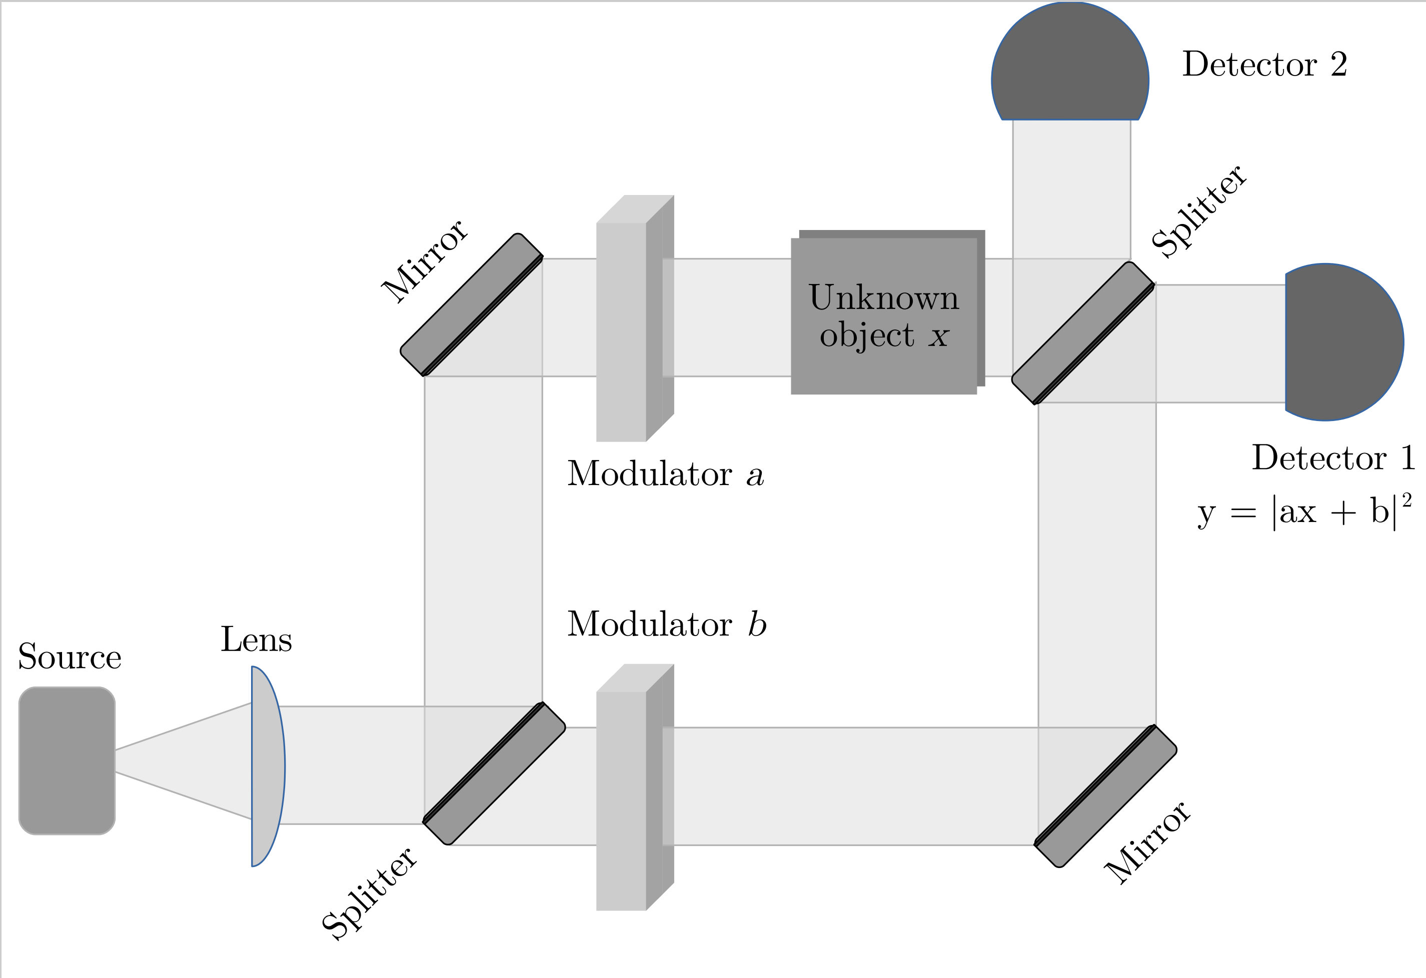

Measurement type:

(Click the radio buttons to change the figure)

- (•) Interferometric measurement
- ( ) Object-free interferometric measurement ($x = 1$)
- ( ) Background intensity measurement ($a = 0$)
- ( ) Transmitted intensity measurement ($b = 0$)

Figure 1: An unknown material sits in one arm of an ideal optical interferometer. The interfering beams are collimated, monochromatic, perfectly coherent, and overlap perfectly at the outputs. The average number of photoelectrons y measured per trial at Detector 1 is $|ax+b|^2$, where a, b, and x are complex numbers. Modulators a and b let us tune the amplitude and phase of a and b respectively. The absorption and phase shift due to the material determines the amplitude and phase of x.

We make a measurement by choosing the amplitude and phase for each modulator and counting photoelectron "clicks" at Detector 1 for some time interval. Suppose that our source and our interferometer are perfectly stable, and our detector never gives false detections[††]. Even in this ideal case, the number of "clicks" we count is inherently stochastic. The number of clicks per trial at Detector 1 is drawn from a Poisson distribution[†††] with expected value $y = |ax+b|^2$. How shall we infer x?

With the probe beam blocked (Fig. 1 with `Measurement type: Background intensity measurement`), we can calibrate the reference beam signal $|b|^2$ due to Modulator b. Since the mean y of a Poisson process grows faster than its standard deviation $\sqrt{y}$, we can calibrate the output amplitude of Modulator b with arbitrarily high signal-to-noise ratio $y/\sqrt{y}$, by chasing the limit $|b|^2 \to \infty$. Next, we can remove the object (Fig. 1 with `Measurement type: Object-free interferometric measurement`) and use a similar procedure to calibrate the output amplitude of Modulator a, and also to calibrate how the modulators tune the relative phase $\angle\frac{a}{b}$ between the probe and reference beams.

With perfect knowledge and control of a and b, we're ready to measure the complex transmission x of our object. Let's add one additional constraint: for this measurement, we can't use an infinite amount of light. Perhaps we're time-limited; our source has finite brightness, we have many objects to measure today, and we can't spend as much time on one measurement as we did on calibration. Perhaps we're dose-limited; too much light might melt or burn our unknown object. Either way, if the cumulative "light dose" $\sum_i^N |a_i|^2$ at trial N exceeds our "dose budget" A, we must stop measuring, and Poisson noise limits our signal-to-noise ratio. What is our signal, and what is our noise?

Let's start with a simple measurement of $|x|$ (how much the object reduces the probe beam's amplitude; Fig. 1 with `Measurement type: Transmitted intensity measurement`). Each trial, Detector 1 yields a number of "clicks" drawn stochastically from a Poisson distribution with expectation value $|ax|^2$. Our noise is the standard deviation of this distribution, $|ax|$. Given our perfect knowledge of a, our mean signal is the *change* in intensity $|a|^2 - |ax|^2$, yielding a signal-to-noise ratio (SNR) of $|\frac{a}{x}| - |ax|$ for a single trial. Note that in the $|x| \to 1$ limit of perfect transmission, the SNR tends to zero. Similarly, in the $x \to 0$ limit of perfect opacity, the SNR tends to infinity! This actually makes sense: an object with sufficiently high transmission becomes invisible to a transmission measurement, and a perfectly opaque object yields a detectable and fluctuation-free signal. Perhaps most interestingly, if you count even a single "click", you know with complete confidence $|x| \neq 0$ and the object was not perfectly opaque.

If we want to measure phase $\angle x$ in addition to amplitude $|x|$, we must use interference (Fig. 1 with `Measurement type: Interferometric measurement`). Our noise in this case is $|ax + b|$, and our signal is still the change in intensity due to the object, $|a + b|^2 - |ax + b|^2$, yielding an SNR of $\frac{|a+b|^2}{|ax+b|} - |ax + b|$. Like transmission measurements, our SNR can be infinite, but since we control a and b, we can achieve infinite SNR regardless of the value of x! All we have to do is choose $\frac{b}{a} = -x$, and we're guaranteed to get a detectable, fluctuation-free signal with infinite SNR: zero "clicks". This is presumably insane, but it raises some fun questions.

In order to achieve "infinite SNR", we must have chosen $\frac{b}{a} = -x$. But if we don't know x yet, how are we supposed to choose a and b? Also, observing zero clicks doesn't guarantee $x = \frac{-b}{a}$; maybe we just didn't wait long enough for a click. Infinite SNR is therefore quite different from infinite precision. However, if we observe even a single click on Detector 1, we know with certainty $x \neq \frac{-b}{a}$. This suggests that if we view measurement as asking a question to the unknown object, the question "What are you?" may be very different from the question "Are you *this*?". This second question has only two possible answers: "maybe" (zero clicks after using up your entire dose budget) and "no" (the first click, at which point you should probably halt the measurement and re-tune your modulators). Clearly there are important clues encoded in how quickly the "no" is delivered.

So, what algorithm should we use for tuning our modulators and drawing conclusions about x? My suspicion is the optimal measurement will use an extremely small fixed value for a (virtually guaranteeing zero or one "clicks" per trial), and choose a new $|b|$ and $\angle b$ after every time the detector clicks, hunting for the informative special case $b = -ax$. I also suspect that the new choice of b should be informed by how long each previous measurement took to produce its first click.

But I'm not going to solve this myself; I'd rather play the game with you. What algorithm would you use? Ideally, express this algorithm as a short Python/Numpy script, ideally in a GitHub repository that I can link to here. Feel free to invite anyone who's interested to play.

What does this teach us about measurement?

- Does "infinite SNR" in a single trial actually indicate a high-precision measurement?
- Do we actually benefit from "hunting" for perfect destructive interference at Detector 1, or would we be better off searching a diverse set of a, b combinations?
- How much better could you do if you also had access to "clicks" from Detector 2?

Notes:

[†]This is an extremely common, extremely general problem. For example, this is the one-voxel limit of nearly the entire field of transmitted light microscopy.

[††] This is consistent with the laws of physics, but an awful lot of work to achieve in practice.

[†††] For a discussion of this assertion, and many fascinating related issues, read the beautifully written second chapter of *The Quantum Challenge: Modern Research on the Foundations of Quantum Mechanics* by Greenstein and Zajonc.

What does this have to do with phase contrast microscopy?

In 1902, Siedentopf and Zsigmondy described the "ultramicroscope" [Siedentopf 1902], an early form of "dark-field" microscopy. By 1920, dark-field microscopy was a mature, well-established technique [Gage 1920], and Zsigmondy was awarded the 1925 Chemistry Nobel “for his demonstration of the heterogenous nature of colloid solutions and for the methods he used, which have since become fundamental...”.

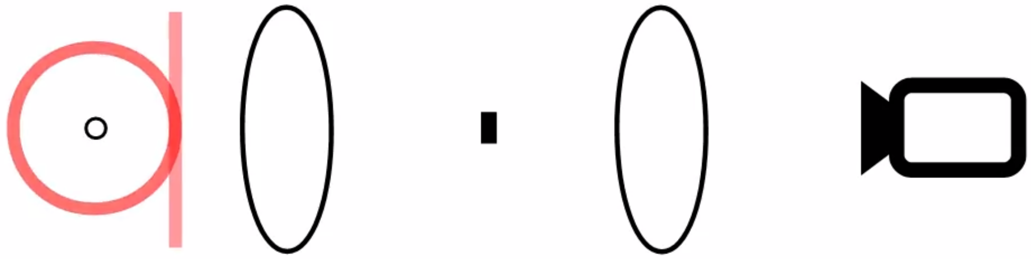

Figure 2: Animation illustrating the key concepts of dark-field microscopy. A red light pulse illuminates a point-like particle (small circle, left). The particle scatters some of this illumination light (expanding red circle). A lens (black oval, left) focuses the illumination pulse onto a beam block (small rectangle, center), which is opaque and absorbs the illumination. Any scattered light that passes through the first lens and is not absorbed by the beam block will pass through a second lens (black oval, right) toward a multi-pixel detector (camera icon, far right), forming an image of the scattering particle on a dark background.

Figure 2 above illustrates the essential concept of dark-field microscopy: only the portion of the illumination beam which is scattered by the sample can reach the detector. Scattering objects in the sample appear as if they are emitting light, and this light is imaged onto the detector on a dark background.

Decades later, Frits Zernike demonstrated "phase contrast" microscopy, [Zernike 1935, Zernike 1942]. Zernike was awarded the 1953 Nobel prize in Physics "for his demonstration of the phase contrast method, especially for his invention of the phase contrast microscope." [Zernike 1955].

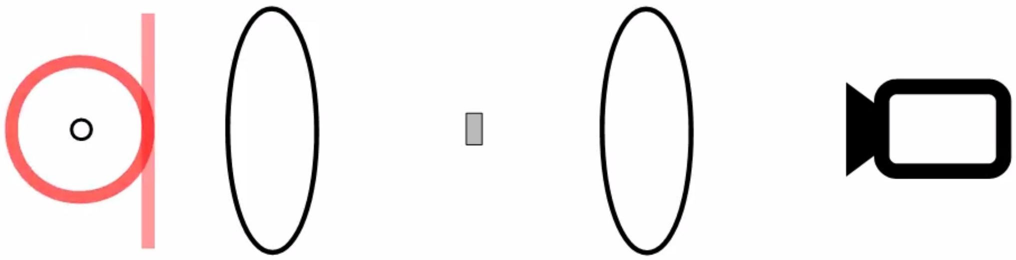

Relative phase between background and scattered light:
(Click the radio buttons to change the figure)

◉ Constructive interference
○ Destructive interference

Figure 3: Animation illustrating the key concepts of phase-contrast microscopy. A red light pulse illuminates a point-like particle (small circle, left). The particle scatters some of this illumination light (expanding red circle). A lens (black oval, left) focuses the illumination pulse onto a beam block (small rectangle, center), which is partially opaque and attenuates the illumination, and also retards its phase. Any scattered light that passes through the first lens and is not absorbed by the beam block will pass through a second lens (black oval, right) toward a multi-pixel detector (camera icon, far right), forming an image of the scattering particle on a bright background.

As illustrated in Figure 3 above, phase contrast microscopy is conceptually similar to dark-field microscopy. To convert our cartoon dark-field microscope to a phase-contrast microscope, we simply replace the opaque beam block with a partially transmissive one. The attenuated illumination beam yields a uniform background on the detector, which interferes coherently with the scattered light image. The partially transmissive beam block also retards the phase of the background, and we can control this retardation (e.g. by tuning the thickness of the block).

For an appropriate retardation, the interference between the scattered light and the uniform background can be primarily constructive (Figure 3 with `Relative phase between background and scattered light: Constructive interference`). In this mode, scattering objects in the sample appear as if they are emitting light, and this light is imaged onto the detector on a bright background.

It may not be immediately obvious why adding a bright background to a Nobel-prize-winning background-free measurement merits another Nobel prize. Certainly, adding background light to an optical measurement adds additional noise [Pawley 2006]. The critical detail is that Zernike added a bright *coherent* background, which adds both additional noise and additional signal. How much signal, and how much noise?

The total intensity I at a pixel of the detector in Figure 3 is proportional to $|\vec{E}_s|^2 + |\vec{E}_b|^2 + 2\vec{E}_s \cdot \vec{E}_b \cos\theta$, where $\vec{E}_s$ is the electric field at that pixel due to light scattered by the point-like particle, $\vec{E}_b$ is the electric field at that pixel due to the coherent background, and θ is the relative phase between the two fields at that pixel. [Fowles 1989]. The number of photoelectrons detected at that pixel during a measurement with duration τ is drawn from a Poisson distribution with mean $\epsilon\tau I$ and standard deviation $\sqrt{\epsilon\tau I}$, where ϵ is a proportionality constant depending on properties of the detector [Fried 1965].

In the case of dark-field microscopy, $\vec{E}_b = 0$. If there is no "stray" light or other noise sources, the presence of the point-like particle changes the mean number of photoelectrons detected by our pixel, from zero to $\epsilon\tau|\vec{E}_s|^2$. The standard deviation in the number of photoelectrons is the square root of this mean. If we call the change in mean photoelectron number due to the point-like particle the "signal", and the standard deviation the "noise", the per-pixel signal-to-noise ratio of dark-field microscopy is $\frac{\epsilon\tau|\vec{E}_s|^2}{\sqrt{\epsilon\tau|\vec{E}_s|^2}} = \sqrt{\epsilon\tau}|\vec{E}_s|$.

In the case of phase contrast microscopy, the presence of the point-like particle changes the mean number of photoelectrons from $\epsilon\tau|\vec{E}_b|^2$ to $\epsilon\tau\left(|\vec{E}_s|^2 + |\vec{E}_b|^2 + 2\vec{E}_s \cdot \vec{E}_b \cos\theta\right)$. If the background $\epsilon\tau|\vec{E}_b|^2$ is extremely stable and well calibrated (as discussed above), and we call the change in mean photoelectron number due to the point-like particle the "signal", and the standard deviation the "noise", the per-pixel signal-to-noise ratio of phase-contrast microscopy is $\frac{\epsilon\tau\left(|\vec{E}_s|^2+|\vec{E}_b|^2+2\vec{E}_s\cdot\vec{E}_b\cos\theta\right)-\epsilon\tau|\vec{E}_b|^2}{\sqrt{\epsilon\tau\left(|\vec{E}_s|^2+|\vec{E}_b|^2+2\vec{E}_s\cdot\vec{E}_b\cos\theta\right)}} = \sqrt{\epsilon\tau}\frac{|\vec{E}_s|^2+2\vec{E}_s\cdot\vec{E}_b\cos\theta}{\sqrt{|\vec{E}_s|^2+|\vec{E}_b|^2+2\vec{E}_s\cdot\vec{E}_b\cos\theta}}$.

In the limits $|\vec{E}_s| \ll |\vec{E}_b|$ (large background), $\vec{E}_s \cdot \vec{E}_b = |\vec{E}_s||\vec{E}_b|$ (parallel fields) and $|\cos\theta| \approx 1$ (retardance tuned for high-contrast interference), we can make the approximations $|\vec{E}_s|^2 \ll 2\vec{E}_s \cdot \vec{E}_b \cos\theta \ll |\vec{E}_b|^2$, and the signal-to-noise ratio of phase-constrast microscopy simplifies considerably to $\sqrt{\epsilon\tau}\frac{2|\vec{E}_s||\vec{E}_b|\cos\theta}{\sqrt{|\vec{E}_b|^2}} = 2\sqrt{\epsilon\tau}|\vec{E}_s|\cos\theta$. In this limit, we see that phase contrast microscopy doubles the signal-to-noise ratio of dark-field!

I suspect that the 1953 Nobel was given for a more practical reason: in addition to doubling the signal-to-noise ratio, phase contrast microscopy (with well-tuned retardance) greatly amplifies both the photon-counting noise *and the signal*. Other sources of noise (e.g. stray light) are "washed out" by the coherently amplified signal. Zernike discusses this point in his Nobel lecture:

> *In my own experiments I could go down to 4-percent transmission, that is, a 5-times enhanced contrast, the limit being set by the unavoidable stray light. It is only under especially favorable circumstances that a higher increase has been attained by the French astronomer Lyot. In his study of the minute ripples of polished lens surfaces he had independently rediscovered phase contrast and could use strips that diminished the amplitude to one-thirtieth, so that ripples only one one-thousandth of a wavelength high showed in good contrast. [Zernike 1955]*

This raises a natural question: in the case of *no* stray light, for a given signal field $\vec{E}_s$, what choice of background field amplitude $|\vec{E}_b|$ and relative phase θ will give optimal signal-to-noise ratio? Naively inspecting our previous equations leads us to a curious result. If we happen to choose a background field that perfectly cancels our signal field, that is, $|\vec{E}_s|^2 + |\vec{E}_b|^2 + 2\vec{E}_s \cdot \vec{E}_b \cos\theta = 0$, then our noise goes to zero, but our signal does not!

This is the origin of the question in our title, and the math, coding, and physics questions that follow. While I have no expectation that "infinite" signal-to-noise ratio allows perfect inference of our sample's optical properties, it's still highly suggestive, and interesting to explore what inference it *does* allow. For mathematical simplicity, I reduced the multi-pixel phase-contrast microscope shown in Figure 3 to the single-pixel interferometer shown in Figure 1, but clearly insight in one case leads directly to insight in the other. If the "infinite" signal-to-noise ratio case actually yields superior inference, one could imagine imaging methods that tune a multi-pixel spatial light modulator in one arm of an interferometer to cancel the fields transmitted by an unknown object in the other arm.

I also think the underlying inference question is a timeless combination of importance and simplicity, which is why I've reduced it to a math puzzle, a coding challenge, and a physics question, rather than simply answering it. I think this document does more good as a challenge than as a reference.

Who wants to play my game?

References

1. [Siedentopf 1902] Uber Sichtbarmachung und Größenbestimmung ultramikoskopischer Teilchen, mit besonderer Anwendung auf Goldrubingläser; H. Siedentopf and R. Zsigmondy; Ann. Phys., 315: 1-39 (1902) https://doi.org/10.1002/andp.19023150102
2. [Gage 1920] Modern dark-field microscopy and the history of its development; S.H. Gage; Transactions of the American Microscopical Society Vol. 39, No. 2 pp. 95-141 Apr (1920) https://doi.org/10.2307/3221838
3. [Zernike 1935] Das Phasenkontrastverfahren bei der mikroskopischen Beobachtung; F. Zernike; Z. technische Physik vol 16, p454–457 (1935) https://doi.org/10.1038/ncb1942
4. [Zernike 1942] Phase contrast, a new method for the microscopic observation of transparent objects; F. Zernike; Physica, Vol 9, p686-698 (1942) https://doi.org/10.1016/S0031-8914(42)80035-X
5. [Zernike 1955] How I discovered phase contrast; F. Zernike; Science 121.3141 (1955): 345-349. https://doi.org/10.1126/science.121.3141.345
6. [Pawley 2006] Handbook of Biological Confocal Microscopy, third edition; J. Pawley; Springer US, ISBN 978-0-387-25921-5, eBook ISBN 978-0-387-45524-2, Chapter 2 (2006) https://doi.org/10.1007/978-0-387-45524-2
7. [Fowles 1989] Introduction to modern optics, second edition; Grant R. Fowles; Courier Corporation, ISBN 0486659577, 9780486659572, Chapter 3 (1989) https://books.google.com/books?id=SL1n9TuJ5YMC
8. [Fried 1965] Noise in Photoemission Current; D. L. Fried; Applied Optics Vol. 4, Issue 1, pp. 79-80 (1965) https://doi.org/10.1364/AO.4.000079
