

This is the accepted manuscript made available via CHORUS. The article has been published as:

Image enhancement in a miniature self-imaging degenerate optical cavity

Albert Ryou, Shane Colburn, and Arka Majumdar

Phys. Rev. A **101**, 013824 — Published 22 January 2020

DOI: [10.1103/PhysRevA.101.013824](https://doi.org/10.1103/PhysRevA.101.013824)

Image enhancement in a miniature self-imaging degenerate optical cavity

Albert Ryou¹, Shane Colburn¹, Arka Majumdar^{1,2}

¹ *Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98195, USA*

² *Department of Physics, University of Washington, Seattle, WA 98195, USA*

(Dated: January 6, 2020)

Image-enhancing optical cavities offer an intriguing platform for nonlinear optical processing of two-dimensional signals, with potential applications in optical information processing and optically implemented artificial neural networks. Here, we analyze the performance of a self-imaging degenerate cavity via numerical simulations and show the existence of a cavity size-dependent minimum spread in the transverse mode resonances. This non-degeneracy, in turn, leads to an inherent trade-off between the resonant enhancement and the fidelity of the intracavity image as the cavity finesse changes.

I. INTRODUCTION

Nonlinear processing of two-dimensional optical signals, or images, holds great potential for many research fields, including optical information processing [1], parametric amplification [2], and optical artificial neural networks [3–5]. Generating optical nonlinearities, however, typically requires high-intensity pulsed lasers, limiting the activity to only laboratory settings. An alternative method is to employ an optical cavity to resonantly enhance the intracavity field of a continuous-wave laser. While cavities have long been used for nonlinear optics [6–8], they have generally been operated in single mode, amplifying one specific optical frequency and suppressing the rest of the spectrum. Because the modes also have distinct spatial patterns, a typical cavity acts as both a spectral and a spatial filter. An attempt to inject a monochromatic image into such a cavity will thus result in its decomposition, with only the resonance-matching spatial component being transmitted.

From an optical information processing standpoint, this essentially means that a typical cavity can only operate on a one-dimensional data stream, eliminating the massive inherent parallelism that is the main advantage of free-space optics over electronics and integrated photonics [9]. To simultaneously operate on two-dimensional signals, a cavity must be capable of transmitting and enhancing an image as a whole; in other words, the cavity's transverse modes that comprise the image must be spectrally degenerate. A class of cavities known as degenerate cavities exhibit multiple modes that satisfy the same resonance condition. Using the paraxial approximation that light rays travel at small angles relative to the cavity's optical axis, Arnaud explored theoretically a special subset of degenerate cavities, called self-imaging cavities, for which all transverse modes are completely degenerate [10]. This concept has been the basis of Ref. [11], which experimentally demonstrated image transmission and frequency doubling with high fidelity in a 20-cm-long cavity.

Recent advances in nanophotonics have made it timely and relevant to revisit the use of the self-imaging degenerate cavity in the context of miniaturized optics. However,




FIG. 1. (Color online) **Cavity geometry and mode spectrum.** (a) Schematic diagram of the cavity. The cavity consists of a flat mirror, a thin lens with focal length f , and a curved mirror with radius of curvature R , separated by d_1 and d_2 . The optical elements have a transverse extent defined by the aperture diameter a . (b) Mode spectrum of a cavity with $f = R = d_1/2 = 1$ mm, $a = 300$ μm , and $d_2 = d_1 + \Delta d_2$. Mode degeneracy for this cavity is predicted to occur for $\Delta d_2 = 0$. The operating wavelength λ is 1 μm . The amplitude reflection coefficients of the mirrors are set to 0.96, while the lens exhibits unity transmission; (inset) zoomed-in plot of the first set of peaks. The profile of the fundamental mode (first peak) has been overlaid in (a), where the color represents the field intensity. (c) The transverse profiles of the first four peaks are well-approximated by the Laguerre-Gauss (LG) modes.

when optical elements shrink in size, their functionalities can change in unanticipated ways [12, 13]. In particular, the paraxial approximation can no longer be taken for granted. The effect of the approximation's breakdown on the mode degeneracy, and subsequently on the resonant enhancement and the fidelity of the intracavity




FIG. 2. (Color online) **Mode spectrum vs. cavity size.** Plot of the cavity mode spectrum as a function of the change in $d_2 = d_1 + \Delta d_2 - N\lambda/2$, with Δd_2 on the y-axis and axial mode shift number N on the x-axis. As N increases from zero, moving the cavity to a different axial mode, the transverse modes first become increasingly closely-packed before reaching the minimum spread, after which they begin to spread out again. The colorbar represents the square root of the field intensity. Note that for large cavities, the degeneracy condition is reached at $N = 0$, as shown in Fig. 7 in Appendix B.

image, warrants careful investigation when we plan to miniaturize the degenerate cavity.

In this paper, we analyze a miniature self-imaging degenerate cavity and explore the consequence of the non-degeneracy of the transverse modes. Using Fourier optics to model light propagation, we demonstrate the formation of individual modes as well as the coherent buildup of an intracavity image. We then study the relationships among image enhancement, fidelity, and cavity size. Finally, we investigate the trade-off between enhancement and fidelity resulting from tuning the cavity finesse.

II. CAVITY SIMULATION

The cavity, shown in Fig. 1a, is composed of three optical elements: a flat mirror, a thin converging lens with focal length f , and a curved mirror with radius of curvature R . The distance between the flat mirror and the lens is d_1 , and the distance between the lens and the curved mirror is d_2 . All the optical elements share the same transverse extent denoted by the aperture diameter a . The two mirrors are partially reflective with the amplitude reflection coefficient r , while the lens exhibits unity transmission.

To simulate light propagation inside the cavity, we use an iterative approach based on the Fox and Li method [14] (see Appendix A). The modes of the cavity are found by injecting a plane wave from the side of the flat mirror. Each round trip yields a modified field profile, and the total intensity, given by the absolute square of the sum




FIG. 3. (Color online) **Minimum mode spread vs. cavity size.** Plot of the minimum spread as a function of d_1 , which parametrizes the cavity length as well as f , R , and d_2 . The aperture diameter of the cavity is set to $\sqrt{50}w$, where w is the largest beam width of the fundamental mode inside the cavity. The spread is smaller for larger d_1 , indicating that the cavity becomes more degenerate as its dimensions grow. The curve is a guide to the eyes only.

of the individual fields after each round trip, reaches a steady state after a sufficient number of round trips.

Computing the cavity mode spectrum requires repeating the process and calculating a series of total intensities while tuning one of its parameters, typically either the wavelength of the incident light or the length of the cavity. We chose to tune the cavity length by sweeping d_2 , the distance between the lens and the curved mirror, while leaving all other parameters fixed. Figure 1b shows a plot of the mode spectrum of a cavity with $f = R = d_1/2 = 1$ mm, $a = 300$ μm , and $d_2 = d_1 + \Delta d_2$, where Δd_2 runs from zero to $\lambda = 1$ μm . As expected, the spectrum exhibits two sets of sharp, closely-packed peaks, one for each axial mode of the cavity. The inset shows a zoomed-in view of the first set. Figure 1c shows the intensity profiles of the first four peaks in the set, which resemble the lowest-order Laguerre-Gauss (LG) modes with $l = 0$ and $p = 0, 1, 2$, and 3 [15]. We note that the physical dimensions of the optical elements of our simulated cavity are about fifty times smaller than those of the previously reported experimental cavity [11].

III. ANALYSIS OF THE MODE DEGENERACY

For infinite-aperture paraxial cavities, it has been shown that the Gouy phase that is responsible for the distinct frequencies of the transverse modes becomes zero for several cavity configurations, leading to the completely degenerate spectrum [10] (see Appendix B). For the cavity configuration shown in Fig. 1a, mode degeneracy is predicted to occur when $d_1 = f + f^2/R$ and $d_2 = f + R$,

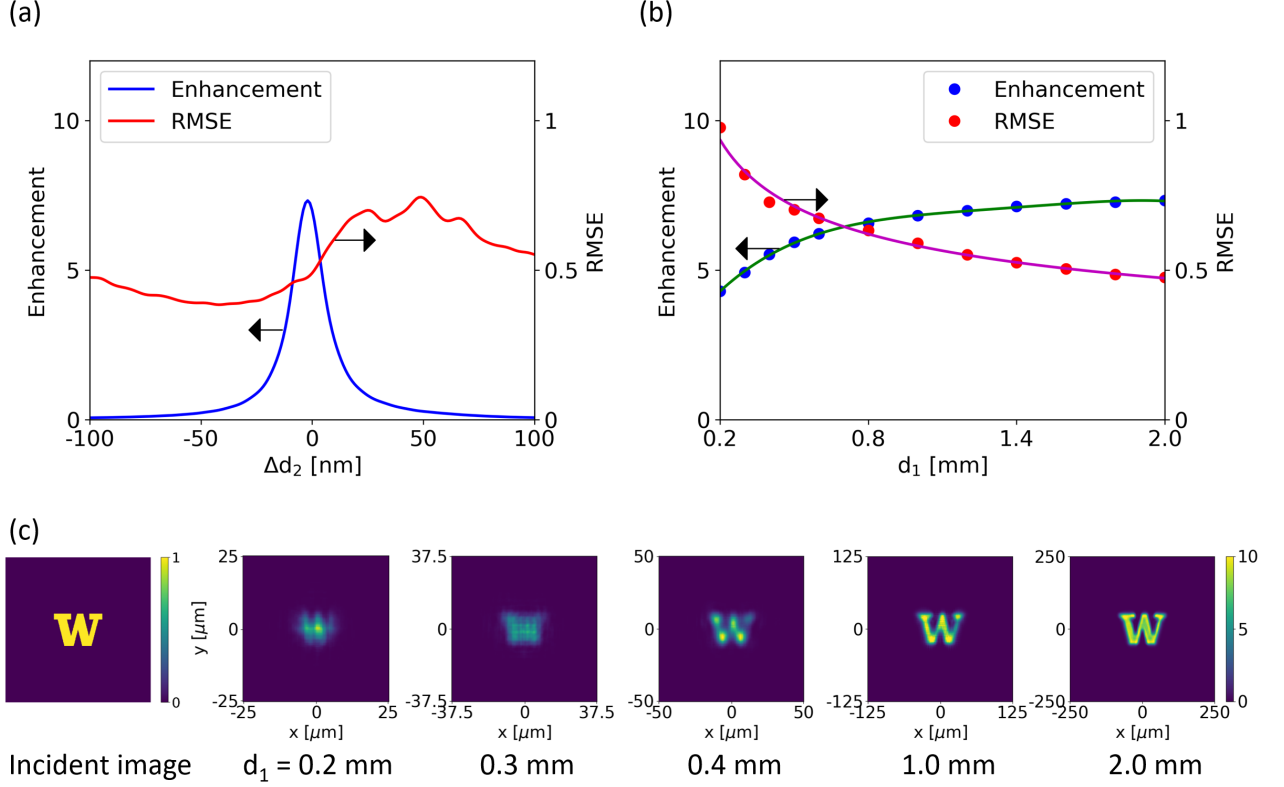


FIG. 4. (Color online) **Injecting an image into the cavity.** (a) Plot of the intracavity field intensity enhancement (blue, left arrow) and RMSE (red, right arrow) vs. Δd_2 . The enhancement is the highest where the cavity modes occur, i.e. $\Delta d_2 \approx 0$. (b) Enhancement (blue, left arrow) and RMSE (red, right arrow) vs. cavity size, parametrized by d_1 . As can be predicted from the result of Fig. 3, as the cavity size decreases, the image becomes both less intense and less similar to the incident image. The error bars are smaller than the points. The curves are guides to the eyes only. (c) Intensity profiles of the incident image, along with the intracavity image for $d_1 = 0.2, 0.3, 0.4, 1.0$, and 2.0 mm.

which becomes, for equal focal length and radius of curvature, $f = R = d_1/2 = d_2/2$. As Fig. 1b shows, however, the resulting simulated modes, while closely-packed, are not degenerate. We attribute this non-degeneracy with miniaturization to the combination of several factors, including the focal shift and the non-paraxial nature of the intracavity light, as elaborated below. A contribution from the numerical effect stemming from the cavity's finite aperture is also described in Appendix C.

It is well-known that a collimated beam incident on an ideal lens focuses slightly inside the lens' geometrical focus f [16]. The deviation, known as the focal shift, is given by $\Delta f = z_R^2/f$, where the Rayleigh range $z_R = \pi w^2/\lambda$, and w is the $1/e^2$ intensity radius of the fundamental cavity mode at a focus [15]. For typical focused beams in large optical systems, $z_R \ll f$ such that Δf is negligible; for our miniature degenerate cavity, on the other hand, both f and z_R near the flat mirror are of the same order of magnitude, resulting in non-zero Δf . This suggests then that we may compensate for the focal shift and reduce the spread among the modes by making the cavity shorter. To test this hypothesis, we shift d_2 away from the predicted degeneracy condi-

tion by setting $d_2 = d_1 + \Delta d_2 - N\lambda/2$, where N is an integer denoting the number of axial modes by which we shift d_2 . Figure 2 shows a plot of the mode spectrum as a function of Δd_2 (y-axis) and N (x-axis) for a cavity with $f = R = d_1/2 = 1$ mm, $r = 0.993$, and $a = 300 \mu\text{m}$. As N increases, the modes become increasingly closely-packed. Eventually, they reach a point of minimum spread before fanning out again.

We attribute the remaining minimum spread to the breakdown of the paraxial approximation. Gaussian optics and the conventional laser resonator theory that employs the Gouy phase are built on the paraxial approximation that an optical wavefront propagates only at a small angle relative to the optical axis [15]. The degree to which the approximation holds thus depends on the modes' beam width w , with first-order corrections for the non-paraxial field scaled by $(\lambda/w)^2$ [17]. The approximation holds well for large cavities, and hence practically perfect degeneracy can be achieved. For our cavity, however, with $\lambda = 1 \mu\text{m}$ and $w < 10 \mu\text{m}$ near the curved mirror, the paraxial approximation breaks down.

Finally, we study how this minimum spread, which has a direct consequence on the intracavity image quality, de-

depends on the size of the cavity. Intuitively, as the cavity shrinks and w becomes smaller, the intracavity field becomes more non-paraxial, causing the modes to spread further. Figure 3 shows the minimal spread, now defined to be the smallest range achievable among the four lowest LG modes by tuning Δd_2 , versus the cavity size, which is parametrized by d_1 . While sweeping d_1 , we set $f = R = d_1/2$, $a = \sqrt{50}w$, where w is the beam width at the plane of the lens, and N is chosen to produce the smallest spread for each simulation. We note that our parametrization leads to the reduction of both the cavity length and the aperture size, since decreasing the length also decreases w . As can be seen in Fig. 3, the minimum spread has an inverse relationship with the cavity length, the effect of which we explore further in the following section.

IV. IMAGE ENHANCEMENT AND FIDELITY VS CAVITY SIZE

The buildup of an image inside a degenerate cavity is the result of the coherent superposition of the individual modes. In a perfect degenerate cavity, as all the modes are resonant, they are enhanced by the same factor, and the image integrity is preserved. The fact that the mode resonances occur at different lengths of our miniature cavity, however, means that a fixed-length cavity will lead to unequal enhancement of the modes, in turn leading to image distortion. A realistic image may consist of hundreds of cavity modes, and it is difficult to predict the amount of distortion in the image as a function of the spread in the modes, especially since the act of perceiving image quality is largely psychological [18]. For this reason, we chose to quantify the quality of the intracavity image by calculating the intensity enhancement and the root-mean-square-error (RMSE) between the intracavity and the incident images. A large RMSE denotes low image fidelity and vice versa.

Intuitively, a cavity whose modes are more closely-packed should form an intracavity image with higher enhancement and fidelity. To test this idea, for a given cavity size, we perform the cavity simulation, but instead of exciting with a plane wave, we inject an image of a “W” with an uniform intensity. The image is placed in the middle of the cavity, and its lateral dimension is set to a third of the aperture diameter a (see Appendix D for the effect of different incident image sizes). Figure 4a shows the enhancement (blue, left arrow) and the RMSE (red, right arrow) of the intracavity image as we tune d_2 near the minimal-spread point found in the previous section. The cavity parameters are $f = R = d_1/2 = 1$ mm, $d_2 = d_1 + \Delta d_2 - N\lambda/2$, $N = 16$, $a = 250$ μm , and $r = 0.95$. As expected, the enhancement is the highest near where the modes occur, i.e. $\Delta d_2 \approx 0$.

We then analyze the performance of the cavity as it shrinks in all three dimensions. We again parametrize the cavity length by d_1 , while keeping $f = R = d_1/2$. The

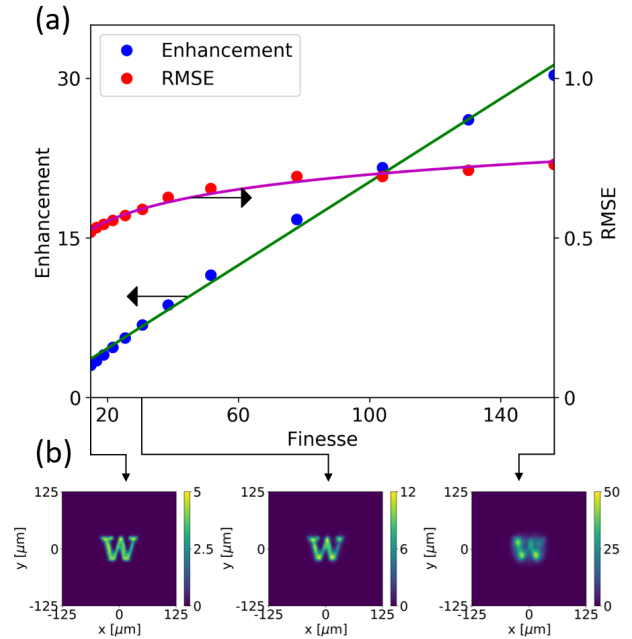


FIG. 5. (Color online) **Intracavity image vs cavity finesse.** (a) Plot of the enhancement (blue, left arrow) and RMSE (red, right arrow) as a function of the cavity finesse. For high mirror reflectivity and high finesse, the modes become more enhanced, but the disparity among the individual modes' enhancement grows. This, in turn, increases the RMSE and distorts the image. On the other hand, for low mirror reflectivity and low finesse, the image is less enhanced but appears more similar to the incident image. The error bars are smaller than the points. The curves are guides to the eyes only. (b) Transverse profiles of the intracavity field intensity for $F \approx 15, 31$, and 160 .

aperture diameter is given by $a = d_1/4$. Thus, both the transverse and the axial dimensions of the cavity maintain their aspect ratio as the cavity size is reduced. The size of the incident image also maintains a fixed ratio relative to a . For each d_1 , we identify N and Δd_2 that yield the highest enhancement and record the values of the enhancement and the RMSE. Figure 4b shows the result of varying d_1 from 0.2 mm to 2 mm. As the cavity shrinks, the enhancement (blue, left arrow) decreases and the RMSE (red, right arrow) increases, which is expected from the cavity size-dependent spread in the modes observed in Fig. 3. Figure 4c shows the intensity profile of the incident image, along with the intracavity intensity for cavities with different d_1 . For $d_1 = 0.2$ mm, the intracavity image is virtually unrecognizable; as d_1 approaches 2.0 mm, it becomes progressively sharper and more similar to the incident image. This constitutes one of the main results of the paper, and it clearly shows how a large cavity can achieve high-fidelity self-imaging.

V. IMAGE ENHANCEMENT AND FIDELITY VS CAVITY FINESSE

While the image enhancement and fidelity are largely determined by the cavity size-dependent spread in the modes, it is possible to increase the effective degeneracy by modifying the cavity finesse. The linewidth Δl of the modes, seen in Fig. 1b, is determined by $\Delta l = \lambda/(2F)$, where the finesse F is a function of the amplitude reflection coefficient r of the mirrors via $F = \pi r/(1 - r^2)$. By decreasing r and therefore F , we can increase Δl , and consequently, the ratio of the enhancements for any pair of modes becomes closer to unity. Thus, although the mode locations remain unchanged, different modes undergo a more equal enhancement, yielding an intracavity image with higher fidelity as a result. Unfortunately, the cost of this effective degeneracy is the reduction in the image enhancement.

Figure 5a shows the enhancement (blue, left arrow) and the RMSE (red, right arrow) as a function of the finesse. Both rise as F increases. The intensity profiles shown in Fig. 5b confirm our expectations. For $F \approx 15$, the intracavity image sharply resembles the incident image but with low enhancement (~ 5); for $F \approx 160$, the enhancement is high (~ 50), but the image appears highly distorted. Thus, depending on the incident image and its modal composition, one can trade off the intracavity enhancement, and hence the efficiency of the nonlinear optical processes, for additional image fidelity. The required level of fidelity in turn may be dictated by specific applications, for instance, the required image classification accuracy for a desired deep learning task.

VI. CONCLUSION

We have simulated a self-imaging degenerate cavity and analyzed the effect of its mode degeneracy on the intracavity image. The results show the critical role that the non-paraxial nature of light plays for image-enhancing cavities even for seemingly macroscopic dimensions on the order of a millimeter. The trade-off among the intracavity image enhancement, fidelity, and cavity size must be carefully considered when employing a self-imaging cavity for specific applications.

One such application of a degenerate cavity is an optical implementation of nonlinear thresholding, one of the

main components for an optical artificial neural network that has never been demonstrated for images [19]. By placing a nonlinear material inside the cavity, such as χ_2 , χ_3 , and two-dimensional materials [20–22], or exploiting self-electro-optic effect [23] or optoelectronic feedback [24], we can engineer pixel-by-pixel interactions between the material and the intracavity field. The output field of the cavity could be read out directly with, for instance, a CCD camera after an optical telescope, or piped into another cavity mimicking a subsequent layer of the neural net. The impact of the nonlinear material on the operation of the degenerate cavity, however, is nontrivial, and warrants further analysis; preliminary results are presented in Appendix E. Finally, we note that in an optical neural network, the input to the cavity is an image transformed by a linear operation, which may or may not appear garbled. Without the loss of generality, here we analyze the performance of the cavity using the image itself. Another possible application is in direct writing of optical tweezing beams [25] or optical potentials [26, 27], which may offer a novel technique for designing new intracavity light-matter interactions.

Finally, building such a cavity with conventional optical elements may be experimentally challenging, due to the required alignment accuracy on the order of a few nanometers, as can be seen from the enhancement peak in Fig. 4a. Multiple precisely-aligned optical elements, on the other hand, can be fabricated by using stacked sub-wavelength diffractive elements such as metasurfaces [28–30], which enable drift- and misalignment-resistant monolithic designs with an exceptionally small footprint. To that end, emerging fabrication technologies such as optical 3D printing can provide attractive solutions [31, 32]. For use in neural networks, the reflectivities of the distributed Bragg reflectors may be chosen to increase either the enhancement or the fidelity at the expense of the other, according to the available training and overall computational resources [5].

VII. ACKNOWLEDGEMENTS

This work was supported by the UW Royalty Research Fund and Air Force grant FA9550-18-1-0104. A.R. acknowledges support from the IC Postdoctoral Research Fellowship and the Mistletoe Research Fellowship. A.M. acknowledges support from the Alfred P. Sloan Research Fellowship.

-
- [1] D. Cotter, R. J. Manning, K. J. Blow, A. D. Ellis, A. E. Kelly, D. Nasset, I. D. Phillips, A. J. Poustie, and D. C. Rogers, Nonlinear optics for high-speed digital information processing, *Science* **286**, 1523 (1999).
 - [2] S. Gigan, L. Lopez, V. Delaubert, N. Treps, C. Fabre, and A. Maître, Continuous-wave phase-sensitive parametric image amplification, *J. Mod. Opt.* **53**, 809 (2006).
 - [3] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljacic, Deep learning with coherent nanophotonic circuits, *Nat. Photon.* **11**, 441 (2017).
 - [4] X. Lin, Y. Rivenson, N. T. Yardimci, M. Velí, Y. Luo, M. Jarrahi, and A. Ozcan, All-optical machine learning using diffractive deep neural networks, *Science* **361**, 1004

- (2018).
- [5] S. Colburn, Y. Chu, E. Shilzerman, and A. Majumdar, Optical frontend for a convolutional neural network, *Appl. Opt.* **58**, 3179 (2019).
 - [6] W. J. Kozlovsky, C. Nabors, and R. L. Byer, Efficient second harmonic generation of a diode-laser-pumped cw nd: Yag laser using monolithic mgo: Linbo₃ external resonant cavities, *IEEE J. Quantum Electron* **24**, 913 (1988).
 - [7] R. W. Boyd, *Nonlinear Optics* (Academic, 1992).
 - [8] T. K. Fryett, K. L. Seyler, J. Zheng, C.-H. Liu, X. Xu, and A. Majumdar, Silicon photonic crystal cavity enhanced second-harmonic generation from monolayer wse₂, *2D Materials* **4**, 015031 (2016).
 - [9] Y. Zuo, B. Li, Y. Zhao, Y. Jiang, Y.-C. Chen, P. Chen, G.-B. Jo, J. Liu, and S. Du, All-optical neural network with nonlinear activation functions, *Optica* **6**, 1132 (2019).
 - [10] J. A. Arnaud, Degenerate optical cavities, *Appl. Opt.* **8**, 189 (1969).
 - [11] B. Chalopin, A. Chiummo, C. Fabre, A. Maître, and N. Treps, Frequency doubling of low power images using a self-imaging cavity, *Opt. Express* **18**, 8033 (2010).
 - [12] K. K. Ghosh, L. D. Burns, E. D. Cocker, A. Nimmerjahn, Y. Ziv, A. El Gamal, and M. J. Schnitzer, Miniaturized integration of a fluorescence microscope, *Nat. Methods* **8**, 871 (2011).
 - [13] K. Liu, S. Sun, A. Majumdar, and V. J. Sorger, Fundamental scaling laws in nanophotonics, *Sci. Rep.* **6**, 37419 (2016).
 - [14] A. G. Fox and T. Li, Resonant modes in a maser interferometer, *Bell Syst. Tech. J.* **40**, 453 (1961).
 - [15] A. E. Siegman, *Lasers* (University Science Books, 1986).
 - [16] S. A. Self, Focusing of spherical gaussian beams, *Appl. Opt.* **22**, 658 (1983).
 - [17] G. P. Agrawal and D. N. Pattanayak, Gaussian beam propagation beyond the paraxial approximation, *J. Opt. Soc. Am.* **69**, 575 (1979).
 - [18] J. B. Phillips and H. Eliasson, *Camera Image Quality Benchmarking* (John Wiley & Sons, 2018).
 - [19] I. A. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, Reprogrammable electro-optic nonlinear activation functions for optical neural networks, *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1 (2020).
 - [20] K. Nozaki, T. Tanabe, A. Shinya, S. Matsuo, T. Sato, H. Taniyama, and M. Notomi, Sub-femtojoule all-optical switching using a photonic-crystal nanocavity, *Nat. Photon.* **4**, 477 (2010).
 - [21] T. K. Fryett, C. M. Dodson, and A. Majumdar, Cavity enhanced nonlinear optics for few photon optical bistability, *Opt. Express* **23**, 16246 (2015).
 - [22] C.-h. Liu, J. Zheng, Y. Chen, T. Fryett, and A. Majumdar, Van der waals materials integrated nanophotonic devices, *Opt. Mater. Express* **9**, 384 (2019).
 - [23] A. Majumdar and A. Rundquist, Cavity-enabled self-electro-optic bistability in silicon photonics, *Opt. Lett.* **39**, 3864 (2014).
 - [24] I. A. D. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, Reprogrammable electro-optic nonlinear activation functions for optical neural networks, *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1 (2020).
 - [25] D. G. Grier, A revolution in optical manipulation, *Nature* **424**, 810 (2003).
 - [26] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold bosonic atoms in optical lattices, *Phys. Rev. Lett.* **81**, 3108 (1998).
 - [27] N. Jia, N. Schine, A. Georgakopoulos, A. Ryou, L. W. Clark, A. Sommer, and J. Simon, A strongly interacting polaritonic quantum dot, *Nat. Phys.* **14**, 550 (2018).
 - [28] N. Yu and F. Capasso, Flat optics with designer metasurfaces, *Nat. Mater.* **13**, 139 (2014).
 - [29] A. Zhan, S. Colburn, C. M. Dodson, and A. Majumdar, Metasurface freeform nanophotonics, *Sci. Rep.* **7**, 1673 (2017).
 - [30] S. Colburn, A. Zhan, and A. Majumdar, Metasurface optics for full-color computational imaging, *Sci. Adv.* **4**, eaar2114 (2018).
 - [31] A. Zhan, R. Gibson, J. Whitehead, E. Smith, J. R. Hendrickson, and A. Majumdar, Controlling three-dimensional optical fields via inverse mie scattering, *Sci. Adv.* **5**, 10.1126/sciadv.aax4769 (2019).
 - [32] T. Gissibl, S. Thiele, A. Herkommer, and H. Giessen, Sub-micrometre accurate free-form optics by three-dimensional printing on single-mode fibres, *Nat. Commun.* **7**, 11763 (2016).
 - [33] J. Goodman, *Introduction to Fourier Optics* (McGraw Hill, 2008).
 - [34] S. Gigan, L. Lopez, N. Treps, A. Maître, and C. Fabre, Image transmission through a stable paraxial cavity, *Phys. Rev. A* **72**, 023804 (2005).
 - [35] E. A. Sziklas and A. Siegman, Mode calculations in unstable resonators with flowing saturable gain. 2: Fast fourier transform method, *Appl. Opt.* **14**, 1874 (1975).
 - [36] K. Matsushima and T. Shimobaba, Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields, *Opt. Express* **17**, 19662 (2009).
 - [37] G. P. Agrawal, Nonlinear fiber optics, in *Nonlinear Science at the Dawn of the 21st Century* (Springer, 2000) pp. 195–211.

Appendix A: Cavity simulation details

1. Light propagation

The simulated optical field is given by a square grid of complex numbers, where the dimension is determined by the aperture diameter a divided by the grid resolution. The grid solution is set to $\lambda/2 = 500$ nm.

The propagation of a field from one plane to another is carried out using the angular spectrum approach, which is based on the fact that any field $u(x, y)$ in real space can be represented as a sum of plane waves with different k vectors [33]:

$$u(x, y) = \iint U(k_x, k_y) e^{i(k_x x + k_y y)} dx dy \quad (\text{A.1})$$

Each plane wave propagates over a distance z , which can be modeled by multiplying the field by the propaga-

tor $H = e^{ik_z z}$, where

$$k_z = \sqrt{k^2 - k_x^2 - k_y^2} \quad (\text{A.2})$$

and $k = 2\pi/\lambda$. Thus, to propagate a field, we take the Fourier transform of the field, multiply it by H , and take the inverse Fourier transform.

2. Optical elements

Each optical element of the cavity, shown in Fig. 1a, can be represented as a phase mask. The phase profile of the flat mirror is unity. The phase profile of the lens is given by

$$\phi_{\text{lens}}(f) = -\frac{\pi(x^2 + y^2)}{f\lambda} \quad (\text{A.3})$$

where x and y denote the transverse coordinates in the simulation grid. The negative sign ensures that the lens is converging. The phase profile of the curved mirror is given by that of a lens with $f = R/2$. The use of the parabolic instead of spherical phase profiles excludes spherical aberration.

Besides the phase profile, the optical elements also have reflectivities that simulate optical loss from either reflection or transmission. For simplicity, the reflectivities for both the flat and the curved mirrors are set equal to r . For the lens, the reflectivity is set to zero so that all light is transmitted without loss.

Thus, the action of an optical element on an incident field is to multiply it by the phase profile and the reflectivity pixel-by-pixel:

$$u'(x, y) = r e^{i\phi(x, y)} u(x, y) \quad (\text{A.4})$$

Additionally, we implement a sharp-edged circular aperture, whose diameter a is equal to the side length of the square grid, by incorporating a circular mask on the plane of each optical element.

Note that associated with the size of the aperture is the mode cut-off loss, which, unlike the cavity leakage loss described above, affects the cavity modes unequally. i.e. greater loss for higher-order modes. The aperture diameter a must be appropriately chosen to minimize the cut-off loss, which may contribute to differential enhancement of an image and reduce its fidelity.

3. Cavity round trip

A simulation begins with an incident field u_{inc} that has unity value everywhere in the two-dimensional grid. The field enters the cavity through the flat mirror, with the transmission amplitude coefficient given by $t = \sqrt{1 - r^2}$. The transmitted wave is then ready to make the first round trip inside the cavity.

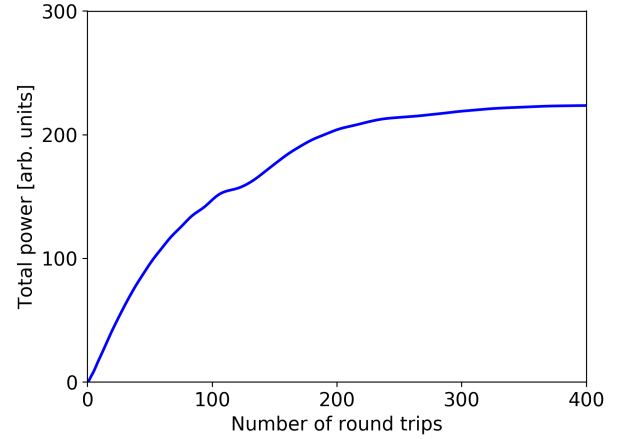


FIG. 6. (Color online) **Intracavity power buildup.** Plot of the intracavity power (sum over all intensity values in the simulation grid at the plane of the flat mirror) vs. the number of cavity round trips. The power increases rapidly and saturates after a sufficient number of round trips, determined by the cavity finesse, or, in terms of the simulation variables, the amplitude reflection coefficient r of the two cavity mirrors.

The round trip consists of a series of alternating free-space propagations and actions by the optical elements. The sequence is as follows: (1) propagation by distance d_1 from the flat mirror to the lens; (2) action of the lens; (3) propagation by distance d_2 from the lens to the curved mirror; (4) action of the curved mirror; (5) propagation by distance d_2 from the curved mirror to the lens; (6) action of the lens; (7) propagation by distance d_1 from the lens to the flat mirror; and finally (8) action of the flat mirror. Together the eight steps comprise one round trip around the cavity.

The total field at the plane of the flat mirror is the sum of the individual fields after making successive round trips. The total intensity is the absolute square of the total field. The total power can be calculated by adding up all the intensity values of the pixels in the grid. We end the simulation when the total power approaches a constant value that indicates that a steady state has been reached between enhancement and loss; see Fig. 6 for the rise and saturation of the total power for a typical cavity mode.

Appendix B: Paraxial degenerate cavity

There are several types of transverse mode degeneracies, depending on the number of modes that share the same Gouy phase. According to paraxial cavity theory, the mode frequencies of a cavity are given by

$$v_{qmn} = \left(q + (s+1) \frac{\alpha}{2\pi} \right) v_{FSR} \quad (\text{B.1})$$




FIG. 7. (Color online) **Paraxial degenerate cavity.** (a) Mode spectrum of a degenerate cavity under the paraxial approximation. $f = R = d_1/2 = 1$ mm, $d_2 = d_1 + \Delta d_2$, $\lambda = 1$ μm , and $r = 0.96$, the same parameters as in Fig. 1b. Unlike the individual peaks that spread out in Fig. 1b, here the modes are highly degenerate, appearing as single peaks at every half wavelength. (b) Zoomed-in spectrum for $\Delta d_2 = 0$. As predicted by paraxial cavity theory, the modes are completely degenerate. The linewidth of the degenerate modes exhibits excellent agreement with the theoretical value given by $\lambda/2F$, where the finesse F is a function of r only. (c) Mode spectrum vs. cavity size. The cavity parameters are the same as those used for Fig. 2. Whereas the length of the non-paraxial cavity in the main text must be offset by an $N \neq 0$ number of half-wavelengths to achieve optimal degeneracy, the paraxial cavity is completely degenerate for $N = 0$, as predicted by the paraxial cavity theory. The colorbar represents the field intensity.

where q denotes the axial mode, and s the transverse mode ($s = 2p + l$ for Laguerre-Gauss modes). The free spectral range is given by $\nu_{FSR} = c/L$, where L is the cavity's round trip distance. The Gouy phase α is related to the eigenvalues of the cavity's round-trip $ABCD$ matrix via:

$$\alpha = \cos^{-1} \left(\frac{A + D}{2} \right) \quad (\text{B.2})$$

where A and D are the matrix's diagonal elements [15].

Different degeneracies can be realized by designing a cavity such that $\alpha = 2\pi K/N$, where K and N are integers [34]. For the self-imaging degenerate cavity, the $ABCD$ matrix is a 2×2 unit matrix when $f = R = d_1/2 = d_2/2$, which leads to $\alpha = 0$. All its transverse modes are completely degenerate, and in principle, any image can be transmitted, and resonantly enhanced via cavity feedback, without distortion. In the geometric optics picture, an arbitrary ray of light in a self-imaging cavity re-traces and returns to the same displacement and slope upon one round trip.

Figure 7 shows the results of the degenerate cavity simulation using the $ABCD$ matrix theory, which is valid under the paraxial approximation. Figure 7a reproduces the mode spectrum as in Fig. 1b, using the same parameters but following the paraxial cavity theory. At the conventional degeneracy condition, the higher-order modes become completely degenerate (Fig. 7b) such that they can be fit with a single Lorentzian curve whose full-width-at-half-maximum agrees with the theoretical value given by $\text{width} = \lambda/(2F)$, where the finesse F is a function of the amplitude reflection coefficient r of the mirrors. Note that for an interferometer whose length is tuned, the linewidth of the mode is independent of the length,




FIG. 8. (Color online) **Mode spectrum vs. aperture size.** Plot of the cavity mode spectrum as a function of the detuning Δd_2 (y-axis) and the aperture diameter (x-axis). As the aperture diameter increases, more peaks are observed in the spectrum, indicating the onset of progressively higher-order modes. At the same time, the modes move closer to one another, even when the size of the aperture starts to become more than an order-of-magnitude bigger than the beam width. The colorbar represents the square root of the field intensity.

as opposed to being inversely proportional to the length when either the frequency or the wavelength is tuned.

Figure 7c reproduces the mode spectrum vs. cavity size plot as in Fig. 2. Again the modes are completely degenerate at the conventional degeneracy condition ($d_1 = d_2, \Delta d_2 = 0, N = 0$), unlike the case for the non-paraxial cavity treated in the main text.

Appendix C: Finite aperture effect

For practical laser cavities with macroscopic mirrors, the aperture can be safely assumed to be infinite. On the other hand, the simulated cavity is necessarily finite in transverse extent due to the limitation in computational resources. Interestingly, the finite aperture results in a numerical effect that appears as a contribution to the spread in the modes. Figure 8 shows the intensity spectrum of a cavity at the predicted degeneracy condition ($f = R = d_1/2 = 1$ mm, $r = 0.99$, $d_2 = d_1 + \Delta d_2$) as a function of Δd_2 (y-axis) and the aperture diameter a (x-axis). As a increases, more modes can “fit” inside the cavity and begin to resonate. In addition, the modes continue to shift and move closer to one another, even when the aperture is an order-of-magnitude larger than the modes.

This effect remains even when we modify the angular spectrum approach with zero-padding the simulation grid, to linearize the discrete Fourier transform [35], and imposing a limit on the field bandwidth [36]. Both methods have been developed to prevent stray light from leaking into “neighboring” cells of the simulation. Further studies are necessary to understand why the size of the simulation grid affects the mode locations, but as can be seen in Fig. 8, the numerical effect is small: a few-nanometer shift in the modes when a changes from 200 μm to 400 μm vs. a tens-of-nanometers shift in the modes for a few-micron tuning of d_2 in Fig. 2.

Appendix D: Effect of the input image size

The image chosen for the simulations is a stylized version of the letter “W”.

As stated in the main text, the enhancement and the fidelity of the intracavity image are heavily dependent on the size of the cavity, which determines the spread in the cavity’s transverse modes. They are also dependent on the modal composition of the incident image itself. The more modes that go into the image’s make up, the greater the distortion in the intracavity image. For the simulations in the main text, the size of the incident image has been fixed to be a third of the grid size. Here we further explore the effect of the size of the incident image on the intracavity image.

Figure 9a shows the enhancement (blue, left arrow) and the RMSE (red, right arrow) of the intracavity field versus the fractional size of the incident image relative to the aperture diameter a . A fraction of 0.2, for instance, indicates that the x-dimension of the incident image is 20% of the x-dimension of $a = 250$ μm . All the other parameters remain fixed during the simulation: $f = R = d_1/2 = 2$ mm and $r = 0.95$. As expected, the enhancement and the RMSE have opposite trends. When the incident image is too small (fraction ~ 0.2), it excites many cavity modes, and as a result, the enhancement is low and the RMSE is high. As the fraction

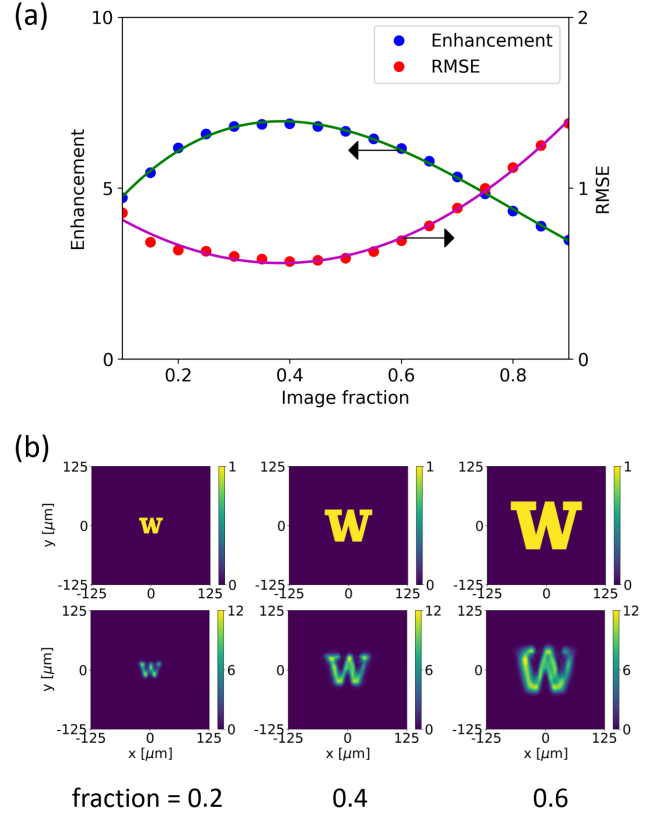


FIG. 9. (Color online) **Intracavity image quality vs incident image size.** (a) Plot of the enhancement (blue, left arrow) and the RMSE (red, right arrow) of the intracavity image as a function of the size of the incident image, as its lateral size is changed from 0.1 to 0.9 of the aperture diameter. The error bars are smaller than the points. The curves are guides to the eyes only. (b) Transverse profiles of the incident image and the intracavity field intensity for different incident image sizes. All three profiles exhibit similar enhancement, but the profile for the fractional size of 0.4 exhibits the highest fidelity. The colorbars indicate the intensity.

increases to about 0.4, the size of the incident image becomes comparable to the beam width of the lowest cavity modes. Thus, it excites only a few modes, yielding the highest enhancement and the lowest RMSE. After hitting this sweet spot, further increasing the incident image size reverses the trends, yielding low enhancement and high RMSE, as the incident image starts to excite higher-order cavity modes again. Figure 9b shows the transverse profiles of the incident image and intracavity field intensities for three different incident image sizes.

Appendix E: Effects of nonlinear media on the cavity mode spectrum

The main objective of our study is the analysis of the mode degeneracy of an empty cavity. On the other hand,

the application of the optical implementation of artificial neural networks requires nonlinear media inside the cavity. A comprehensive study of the impact of nonlinear optical effects on the operation of a near-degenerate cavity is beyond the scope of our paper. Here we discuss two nonlinear effects, saturable absorption and the optical Kerr effect, and outline how they may affect the cavity mode spectrum.

Saturable absorption, or the decrease of light absorption with increasing intensity, occurs when the rate of decay of the medium's excited state becomes comparable to the incident photon flux, leading to the depletion of the ground state. As a result, the input-output behavior is described by [7]

$$I_{out} = I_{in}e^{-A} \quad (\text{E.1})$$

$$A = \frac{\alpha}{1 + I_{in}/I_{sat}} \quad (\text{E.2})$$

where we parametrize the absorption A by α and I_{sat} , which depend on particular material properties.

To implement saturable absorption in our model, we place an infinitesimally thin nonlinear layer at the location of the flat lens. On every round-trip, this layer locally attenuates the intracavity field as a function of the intensity. Figure 10a shows the cavity mode spectrum as in Fig. 1b, but with different values of α . For simplicity, we set $I_{sat} = 1$. As α increases from zero to 0.01, the peaks decrease in height but do not shift. This is reasonable, since saturable absorption only affects the field magnitude and not its phase. For image transmission, the attenuation is not uniform, affecting the brightest mode-component the least. While this may move the location of the highest enhancement (Δd_2 in Fig. 4a), we expect the trade-off on the cavity's enhancement, fidelity, and finesse to remain consistent.

The analysis in the case of the Kerr effect is more complex. The change in the refractive index of a material in response to the intracavity field intensity is described by [7]

$$n = n_0 + n_2 I \quad (\text{E.3})$$

To implement the Kerr effect in our model, we place a $1\text{-}\mu\text{m}$ -thick nonlinear material right after the flat mirror and employ the split-step Fourier method: we alternate between propagating the wavefront $u(x, y)$ via the angular spectrum method and multiplying it by $e^{i\phi}$, where $\phi = n_2|u(x, y)|^2 kl$ [37], $k = 2\pi/\lambda$, and $l = 1\text{ }\mu\text{m}$. The large values of n_2 are chosen to compensate for small l , which greatly speeds up our simulation.

Figure 10b shows the cavity mode spectrum as in Fig. 1b, but with different values of n_2 . As n_2 increases from zero to $0.1\text{ m}^2/\text{W}$, the peaks shift as well as change in shape away from standard Lorentzian curves. The shift in the locations of the peaks can be attributed in part to the thin lens effect [33]. For large n_2 , the peaks begin to exhibit definitively non-Lorentzian shapes, whereby we

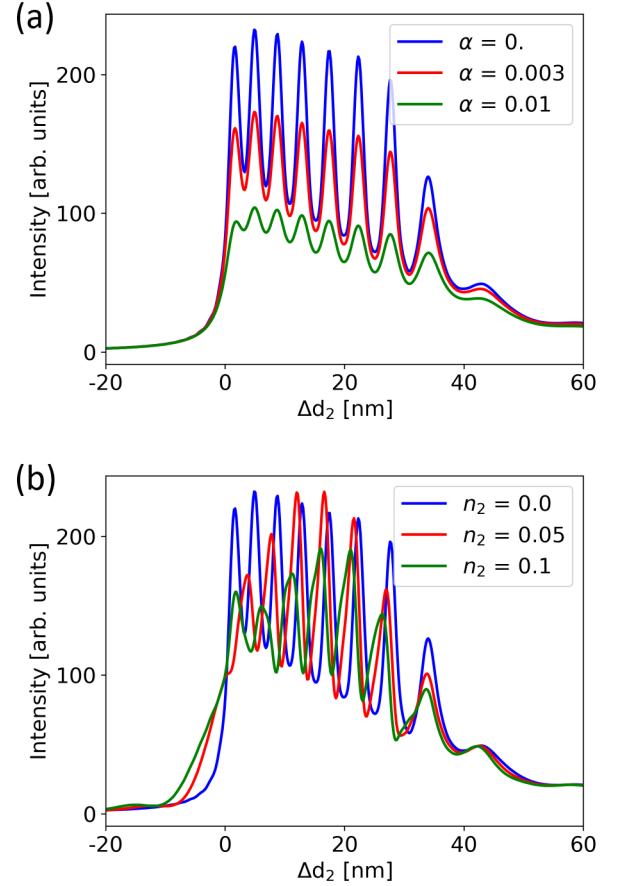


FIG. 10. (Color online) **Cavity mode spectrum with nonlinear media.** (a) Mode spectrum of a cavity ($f = R = d_1/2 = 1\text{ mm}$, $a = 300\text{ }\mu\text{m}$, $d_2 = d_1 + \Delta d_2$, $\lambda = 1\text{ }\mu\text{m}$, $r = 0.993$) containing a saturable absorber. As α increases, the peaks become attenuated, but their locations remain the same. (b) Mode spectrum of a cavity (same geometry as (a)) containing a Kerr medium. As n_2 increases, the peaks become attenuated and shift as well as changing shapes.

conclude that the cavity field can no longer be described in terms of its orthogonal modes. Figure 10b does suggest, however, that up to some n_2 , the total power contained in the image remains concentrated within narrow Δd_2 . In this case, we hypothesize that there still exists some measure of degeneracy that translates to the preservation of an incident image fidelity.

Finally, concerning the optical implementation of the artificial neural networks, we believe that saturable absorption can be implemented, for example, with a vapor cell containing thermal atoms, with the main conclusions of the paper unchanged. The impact of the Kerr effect is more subtle, but since the nature of the neural network allows adaptation (training) of the system based on the output, the nonlinear mixing of the mode components by the activation function itself may prove acceptable.