PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Blending physics with artificial intelligence

Kadambi, Achuta

Achuta Kadambi, "Blending physics with artificial intelligence," Proc. SPIE 11396, Computational Imaging V, 113960B (24 April 2020); doi: 10.1117/12.2565099

Event: SPIE Defense + Commercial Sensing, 2020, Online Only, California, United States

Blending Physics with Artificial Intelligence

Achuta Kadambi[a,b]

[a]Visual Machines Group, UCLA, 420 Westwood Plaza, Los Angeles CA, USA
[b]Department of Electrical Engineering, UCLA, 420 Westwood Plaza, Los Angeles CA, USA

ABSTRACT

For centuries, humans have discovered the physical laws that underpin our world. What if the next Einstein or Newton is not a human, but a machine? Machines that are physics-aware can transform a multitude of fields, poised to enable unexpected and meaningful feats in science and engineering. In this paper, we survey methods germane to the imaging sciences where we observe a very special convergence of a millennia of optical theories with decades of digital photos.

Keywords: Computer Vision, Computational Imaging, Physics-based Vision, Scientific Discovery

1. INTRODUCTION

Imagine if it was possible to teach machines to discover the laws of physics. Such a physics-aware artificial intelligence (AI) can transform a multitude of fields, including the imaging sciences. Smart cameras could enable cars to avoid collisions before they happen. A smart microscope can capture videos of patients and gain intuition about the underlying biophysics of a pathology. Advanced algorithms may even discover new optical laws that have eluded human understanding.

However, in order to realize physics-aware machines, we must work together in advancing progress in computer vision. Although computer vision is one of the most impactful fields of artificial intelligence, the modern paradigm of vision is not physics-aware. Existing algorithms largely tune knobs to parameters. Judea Pearl, a UCLA colleague and Turing Award Laureate has been quoted as saying that "the impressive achievements of deep learning amount to just curve fitting". These algorithms, curve fitting perhaps, then sit on top of ordinary cameras, which capture a miniscule, two-dimensional projection of the trillions of light paths in a scene. As a result, computer vision algorithms lack the sensitivity to detect subtle bio-signals, are easily fooled by spurious reflections and material effects, and are easily fooled by adversarial attacks.

In order to overcome this paradigm, we must incorporate physics into the pipeline of AI, from the processing of AI to the capture of sensory inputs for AI. The author is a member of the Visual Machines Group at UCLA, where the aim is to combine physical insights and computer algorithms to transform imaging in ways that are unexpected, yet meaningful. The group has a unique approach to computer vision, that extends all the way down to the hardware and up to the algorithms. In what follows, we provide a survey of research in two areas: physics-based vision and physics-based AI. We also provide an overview of nascent applications for imaging, ranging from patient health monitoring to new types of long-range camera systems.

2. PHYSICS-BASED ARTIFICIAL INTELLIGENCE

We can divide research in physics-based artificial intelligence into two broad themes, depending on whether the physics is known or unknown. The first theme assumes that physical laws are known, and aims to **blend** the laws of physics with artificial intelligence. In contrast, the second theme assumes that the physics is unknown and seeks to **discover** the laws of physics.

Further author information: (Send correspondence to A.K.)
A.K.: E-mail: achuta@ucla.edu

Computational Imaging V, edited by Lei Tian, Jonathan C. Petruccelli, Chrysanthe Preza, Proc. of SPIE Vol. 11396, 113960B · © 2020 SPIE
CCC code: 0277-786X/20/$21 · doi: 10.1117/12.2565099

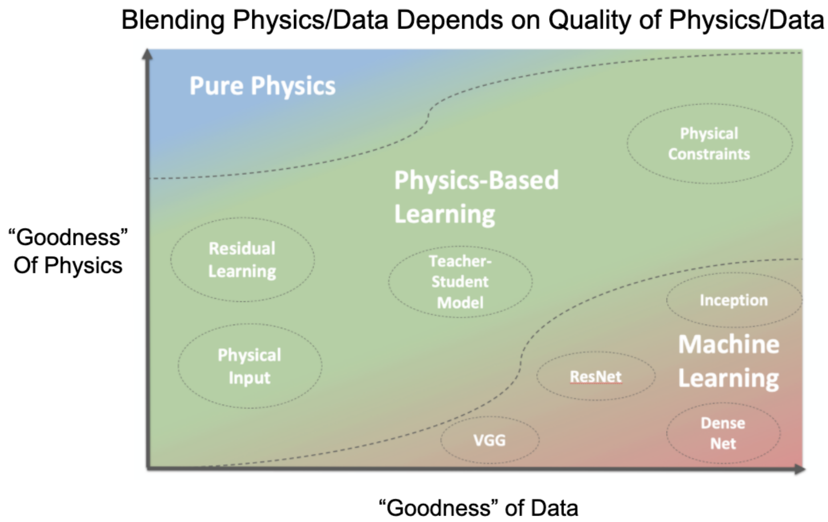

Figure 1. **Blending physics and learning often depends on the quality of physics and data.** The "goodness" of physics and data can be defined more precisely in context of training data or model mismatch. Figure from Ba et al.[1]

2.1 Blending Physics and Learning

Imagine if we are tossing a basketball and would like to estimate the trajectory—will it go in the hoop or land short? Elementary physics can help us solve this problem, but our calculation may be off due to nuisance factors like wind speed or backspin. If we have a multitude of training examples of basketball tosses, we can apply a machine learning model to estimate the trajectory. Blending physics and learning is not new and dates back as early to the Kalman filter.[2] However, we are now at the stage where deep learning models perform much better than classical control algorithms for many problems.

To bridge this gap, the area of physics-based learning (PBL) aims to blend physical priors with modern methods in machine learning, and in particular deep learning, to combine the best of both worlds. The subtlety in these approaches of blending is to understand how to perform the blending in different regimes of physics and data quality. As illustrated in Figure 1, if the physical model is very "good", then a purely physical solution could be used for the blending (e.g. the kinematic equations in the projectile example). In contrast, if the data is "good" then a purely data-driven solution could be used for blending. Here, "goodness" is left open to intepretation—it is often rigorously quantified in some manner, be it the amount of training data or a bound on the degree of physical mimatch. The subtlety occurs in the green region of Figure 1, where data and physics are both informative, but not perfect.

PBL architectures have achieved competitive performance with respect to naive neural networks, on a wide variety of tasks in fields as diverse as computational microscopy,[3–6] low level and high level computer vision,[1,7–9] and medical imaging.[10,11] Unfortunately, these PBL methods are typically designed for a specific task, or regime of Figure 1. Generalization would (as a first step) require a PBL architecture capable of adapting to variations in the correctness of physics or the quality of training data. Experiments show that no such architecture exists, and it is possible to approach this from a different angle as shown in Figure 2. Inspired by work in neural architecture search (NAS),[12–15] Ba and Zhao et al. propose an attempt to automatically find the optimal PBL architecture through a discovery algorithm known as AutoPhysics.[16] This is an important step in adopting physics-based learning to encompass the wide range of physical problems, where priors are only approximate and training data can be sparse.

2.2 Discovering New Laws of Physics

The second theme is discovery, where the goal is to learn a physical model whose form and parameters are unknown. The apocryphal story of how Newton discovered the laws of gravitation seeing an apple is fall can be

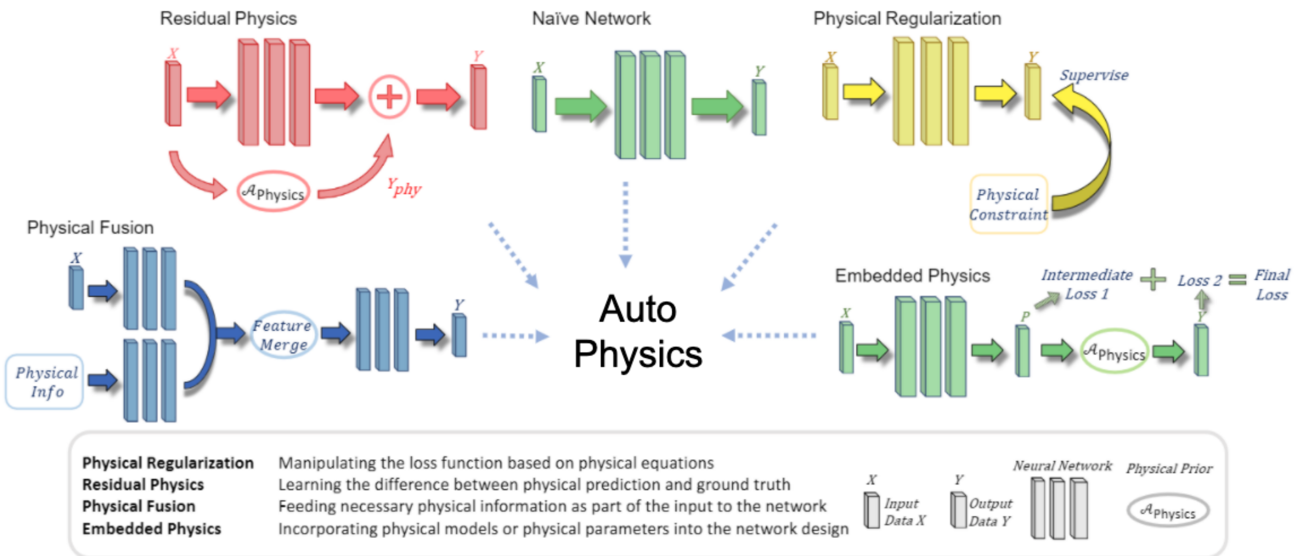

Figure 2. **A new method enables automatic selection of the blending between physics and learning.** Ba and Zhao et al. design a physics-based variant of neural architecture search to find the best performing manner in which to blend physics and learning. Figure from Ba and Zhao et al.[16]

transposed as a computer vision problem. Concretely, imagine a smart camera that observes a dynamic event (e.g. the apple falling) and is able to deduce the physical laws (e.g. gravitation and Kinematics). If the discovery problem can be scaled, it could perhaps aid in discovery of more complex physical models, such as scattering phase functions, reflections distributions, and more.

Progress in discovering physical laws is still nascent. Most previous work achieves success akin to partial discovery,[17–23] where algorithms are either able to discover governing equations or the parameters, but not both. Recent work from Chari et al. takes a closer step toward full discovery.[24] Here, the work takes only a video sequence of bounding boxes as input (e.g. positional data) and is able to learn both the governing equation and rest of the parameters (e.g. velocity and the gravitational constant). Figure shows a few elementary scenes from their paper, where a network takes as input only video sequences and can discover projectile kinematics (linear motion equations) or circular rotation (sinusoidal motion equations). However, these basic scenes are only a starting point. The next frontier will be to scale this to more complex scenarios, and eventually, scale to the unknown physical challenges.

3. PHYSICS-BASED VISION

Physics-based thinking can transform subfields of AI, such as computer vision. Ordinarily, computer vision teaches machines how to see. The conventional method is to feed two-dimensional imagery $I(x, y)$ into neural networks, where x and y are spatial coordinates, either in pixels or meters.

Physics-based vision is a field that draws upon knowledge about the physics of how light forms an image to solve long standing problems in computer vision, such as 3D reconstruction, object recognition, segmentation and more. Progress in this area has been accelerated by thinking beyond the typical 2-dimensional matrix image structure $I(x, y)$ used in computer vision and image processing. One way to express the higher-dimensional relationship is in integral form, where

$$I(x, y) = \int\int\int\int\int I(x, y, \theta_1, \theta_2, \lambda, \rho, t)\, d\theta_1\, d\theta_2\, d\lambda\, d\rho\, dt \qquad (1)$$

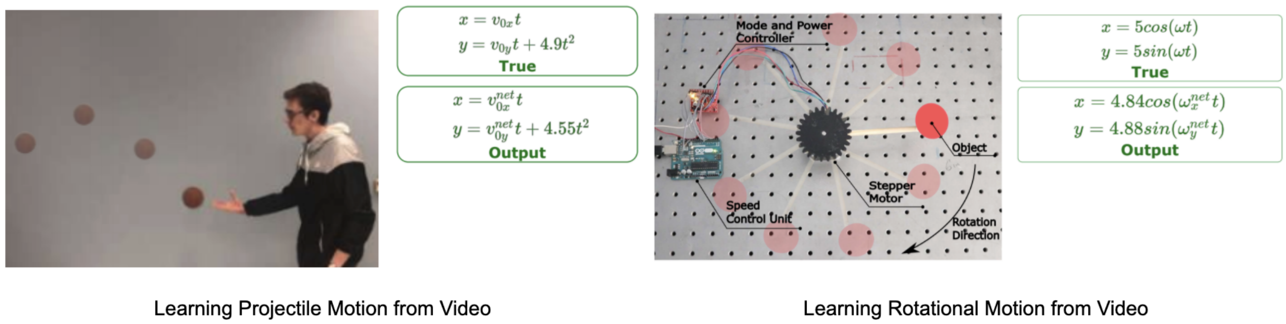

Figure 3. **Discovering physical laws from video streams.** Work by Chari, Talegaonkar, Ba et al. introduce a way to learn governing equations (e.g. projectile motion) and parameters (e.g. velocities) from video streams. The method succeeds on elementary scenes, and it will be exciting to see how such techniques scale to complex physical phenomena. Figure from Chari et al.[24]

represents the low-dimensional projection $I(x, y)$ used in ordinary computer vision and the set $\{\theta_1, \theta_2, \lambda, \rho, t\}$ denotes additional parameters that are not sampled in conventional vision.[25] These parameters include the two light field angles, wavelength, polarization, and time, respectively. While an ordinary camera simply integrates out diversity across angle or polarization, it is possible to modify the imaging setup to capture these dimensions. Then, by sampling these additional dimensions of light, the art of physics-based vision involves linking higher-dimensional image formation models (i.e. higher in dimension than 2D) to key tasks in computer vision. In what follows, we review some work from the computational imaging community in probing specific dimensions of light transport.

3.1 Probing the Temporal Dimension

Whereas an ordinary cameras simply counts the number of photons that strike a sensor pixel, research in **transient imaging** seeks to understand also the timing of photons. The timing characteristics of photons can be used in a classical sense, to obtain 3D shape (e.g. by using the time of flight principle). However, it can also be used for meta-characteristics, such as transient imaging,[26–30] the inference of material properties,[31–33] multipath analysis,[27,28,34–37] and even the ability to see around corners.[38–44]

3.2 Probing the Polarimetric Dimension

Ordinarily, the camera model that is used in computer vision is polarization agnostic. In analogous fashion to the previous subsection, recent work in computational imaging has leveraged the polarization of light to realize extraordinary new capabilities. For example, the work of "Polarized 3D" by Kadambi et al.[45,46] studies how polarization cues can be used to upsample 3D images. The high-quality geometry, at nearly micron-scale is shown in Figure 4 on the left-hand-side. More recently, Kalra et al.[47] also leveraged polarization cues to tackle transparent object instance segmentation. Transparent object segmentation is a hard problem because, in an RGB image, the texture transmitted through the transparent object can overshadow the texture of the transparent object itself. In polarized imagery, the amount of polarized light is measured. Since transparent objects tend to heavily polarize light, they are able to observe the texture of the object itself. Leveraging this, along with their custom deep learning framework, they are able to achieve segmentation robust to print-out spoofs, novel environments, and use this for robotic bin picking of transparent objects. Polarization has also been widely used for imaging through scattering media,[48,49] surface normal reconstruction,[50–53] face and reflectance capture,[54,55] underwater imaging,[56] and even in combination with time of flight imaging.[57] Recently, Tanaka et al. introduced, for the first time, the use of polarization to see around corners.[58]

3.3 Spectral

Leveraging the power of spectral information has been an important theme in physics-based vision. We highlight a few examples germane to computatoinal imaging. Maeda et al.[43] leverage the specular reflectance chracteristics

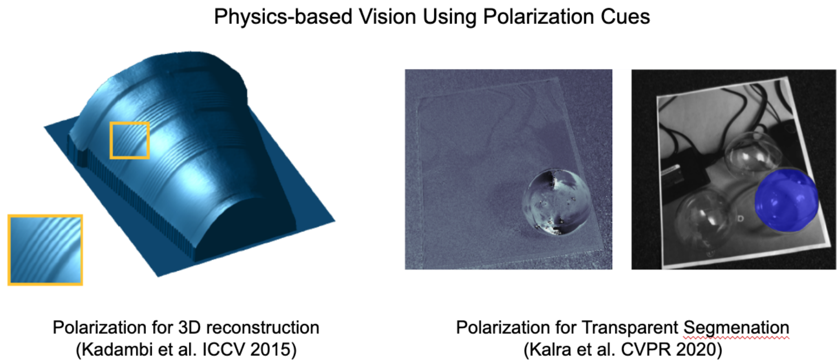

Figure 4. **Harnessing the power of polarization cues to transform computer vision.** At left is a geometric reconstruction of a coffee cup at micron scale using polarization cues to upsample a depth map. At right is the use of polarization cues to segment transparent objects while avoiding print-out attacks (where a picture of an object is used to fool a vision system). Figures from Kadambi et al.[45] and Kalra et al.[47]

of long-wave infrared to see around corners. The strong specular behavior allows for drastic improvements in shape resolution and reconstruction over existing visual methods. For the first time, the authors are able to show reconstruction of a human silhouette around the corner. Along this theme, and although it does not use light, it is worth mentioning the creative use of acoustics to achieve a similar effect of specular reflection by Lindell et al.[59] Other authors like Tanaka et al.,[33] have introduced a novel time-resolved decomposition technique for far infrared light transport, that exploits transience at the much slower speed of propagation of heat. This enables certain transient imaging effects to be observed and analysed at video frame rates. Spectral imaging can also be combined with programmable illumination, as inspired by Saragadam et al.[60] who extend optical computing techniques (analogous to[61]) across the spectral band.

4. ENABLING EXTREME IMAGING USING PHYSICS AND AI

Putting physics and AI together provides a toolbox that can be used to transform imaging into extreme domains. Here we discuss two directions of interest. The first pertains to the design of cameras for satellite imagers, and the second for non-contact health diagnostics. A recurring theme is end-to-end learning of camera systems and low level parameters, analogous to previous work in computational imaging.[62–65]

4.1 Aerial Imaging

Today, satellite imagery is an important tool for many industries, including military, agriculture, construction and real-estate. More-than-ever, the demand for precise, reliable and accessible satellite imagery is required. Continuously pushing the boundary of satellite imaging, such as pushing the resolution limits to 0.1m, will spike new applications in healthcare and computer vision.

However, the traditional design of satellite imaging systems strongly limits its current imaging capacity. Typical satellite imaging systems are monolithic, with constellations acting as independent agents rather than collaborative meta-sensors. Furthermore, optical hardware design choices are heuristic and not optimized for image reconstruction and post-processing. Consequently, bulky and expensive systems are required to achieve high-quality imaging with computationally intensive post-processing techniques, all while being inflexible to adapt to new demands and physical scenarios.

To advance satellite imagery, our group is interested in looking at new forms of satellite perception hardware with data extraction and reconstruction software (Figure 5). Leveraging advances in machine learning and auto-differentiation, we enhance our design space to include the formation of satellites, sensor parameters, laws of physics and active human inputs. By optimizing our hardware and software with user-defined application constraints in an end-to-end manner, we hope to construct a cost-effective, adaptive and intelligent satellite imaging system.

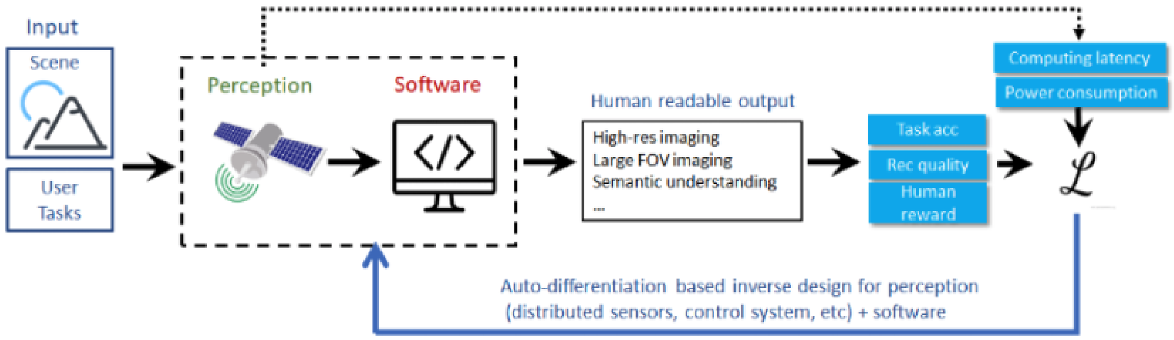

Figure 5. **An example of end-to-end imaging system design for imaging on satellites.** In contrast to studying the imaging system as an isolated unit, the entire input from satellite flight path to imaging focus point can be seen as parameters. Autodifferentiation can learn the optimal combination of joint optical/aeronautical parameters.

4.2 Health

For population-scale screening, a single camera, placed in a busy area of a hospital, can photograph upwards of 100,000 people a month. Although authors like Poh et al.[66] have successfully leveraged cameras for vital isgn measurements, ordinary cameras are somewhat limited. They cannot measure fever, for example, since they do not detect temperature. It turns out that fever (elevated temperature) is an early indicator of illnesses like cancer or infection. For instance, In the currently ongoing COVID-19 epidemic, hospitals do not have the throughput to detect fevers of patients or medical staff because they were using low-throughput instruments to measure temperature. One such device is the non-contact infrared (IR) thermometer, which works by measuring the IR radiation emitted from a patient. Although accurate, these thermometers do not have the necessary throughput to keep up with patient intake. A laser dot must be carefully aligned with the patient's forehead as the instrument outputs a single-point measurement. It is impossible to measure the temperature of multiple individuals with such point-wise devices. Inspired by work from Nowara and colleagues,[67,68] perhaps it may be possible to one day have a thermal camera that can image a swath of people for triaging.

4.3 CONCLUSION

We believe the future of imaging will rely on practitioners who can jointly push the boundaries of artificial intelligence and physics. A snapshot of work covered in this short report is not representative of the richness of topics that can be studied in this umbrella. Focused review articles offer much more detail, and we point the reader to Raghu et al.[69] for a survey on scientific discovery, Willard et al.[70] for a survey on physics-based learning, Maeda et al.[71] for a survey on non-line-of-sight imaging, and Bhandari et al.[72] for a survey on time of flight imaging. The author is currently writing a textbook on computational imaging, to be published by MIT Press in 2020.

ACKNOWLEDGMENTS

Visual Machines graduate students Yunhao Ba, Pradyumna Chari, Chinmay Talegaonkar and Guangyuan Zhao are first authors or co-first authors on many of the papers we review here. We also thank Agastya Kalra (Akasha Imaging) for contributions to the polarization section. We thank Bahram Jalali, Vwani Roychowdhury, and Aydogan Ozcan for fruitful discussions.

REFERENCES

[1] Ba, Y., Chen, R., Wang, Y., Yan, L., Shi, B., and Kadambi, A., "Physics-based neural networks for shape from polarization," *arXiv preprint arXiv:1903.10210* (2019).

[2] Kalman, R. E., "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering* **82**(Series D), 35–45 (1960).

[3] Rivenson, Y., Wu, Y., and Ozcan, A., "Deep learning in holography and coherent imaging," *Light: Science & Applications* **8**(1), 1–8 (2019).

[4] Nehme, E., Weiss, L. E., Michaeli, T., and Shechtman, Y., "Deep-storm: super-resolution single-molecule microscopy by deep learning," *Optica* **5**(4), 458–464 (2018).

[5] Nguyen, T., Xue, Y., Li, Y., Tian, L., and Nehmetallah, G., "Deep learning approach for fourier ptychography microscopy," *Optics express* **26**(20), 26470–26484 (2018).

[6] Sinha, A., Lee, J., Li, S., and Barbastathis, G., "Lensless computational imaging through deep learning," *Optica* **4**(9), 1117–1125 (2017).

[7] Goy, A., Arthur, K., Li, S., and Barbastathis, G., "Low photon count phase retrieval using deep learning," *Physical review letters* **121**(24), 243902 (2018).

[8] Metzler, C. A., Schniter, P., Veeraraghavan, A., and Baraniuk, R. G., "prdeep: Robust phase retrieval with a flexible deep network," *arXiv preprint arXiv:1803.00212* (2018).

[9] Sun, Y., Liu, J., and Kamilov, U. S., "Block coordinate regularization by denoising," *arXiv preprint arXiv:1905.05113* (2019).

[10] Jin, K. H., McCann, M. T., Froustey, E., and Unser, M., "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing* **26**(9), 4509–4522 (2017).

[11] Kang, E., Min, J., and Ye, J. C., "A deep convolutional neural network using directional wavelets for low-dose x-ray ct reconstruction," *Medical physics* **44**(10), e360–e375 (2017).

[12] Zoph, B. and Le, Q. V., "Neural architecture search with reinforcement learning," *arXiv preprint arXiv:1611.01578* (2016).

[13] Baker, B., Gupta, O., Naik, N., and Raskar, R., "Designing neural network architectures using reinforcement learning," *arXiv preprint arXiv:1611.02167* (2016).

[14] Liu, H., Simonyan, K., and Yang, Y., "Darts: Differentiable architecture search," *arXiv preprint arXiv:1806.09055* (2018).

[15] Cai, H., Zhu, L., and Han, S., "Proxylessnas: Direct neural architecture search on target task and hardware," *arXiv preprint arXiv:1812.00332* (2018).

[16] Ba, Y., Zhao, G., and Kadambi, A., "Blending diverse physical priors with neural networks," *arXiv preprint arXiv:1910.00201* (2019).

[17] Huang, S., Cheng, Z.-Q., Li, X., Wu, X., Zhang, Z., and Hauptmann, A., "Perceiving physical equation by observing visual scenarios," *arXiv preprint arXiv:1811.12238* (2018).

[18] Wu, J., Yildirim, I., Lim, J. J., Freeman, B., and Tenenbaum, J., "Galileo: Perceiving physical object properties by integrating a physics engine with deep learning," *NeurIPS* (2015).

[19] Brubaker, M. A., Sigal, L., and Fleet, D. J., "Estimating contact dynamics," *ICCV* (2009).

[20] Bhat, K. S., Seitz, S. M., and Popovic, J., "Computing the physical parameters of rigid-body motion from video," *ECCV* (2002).

[21] Mottaghi, R., Bagherinezhad, H., Rastegari, M., and Farhadi, A., "Newtonian image understanding: Unfolding the dynamics of objects in static images," *CoRR* **abs/1510.04048** (2015).

[22] Purushwalkam, S., Gupta, A., Kaufman, D. M., and Russell, B., "Bounce and learn: Modeling scene dynamics with real-world bounces," *arXiv preprint arXiv:1904.06827* (2019).

[23] Wu, J., Lu, E., Kohli, P., Freeman, B., and Tenenbaum, J., "Learning to see physics via visual de-animation," *NeurIPS* (2017).

[24] Chari, P., Talegaonkar, C., Ba, Y., and Kadambi, A., "Visual physics: Discovering physical laws from videos," *arXiv preprint arXiv:1911.11893* (2019).

[25] Adelson, E. H., Bergen, J. R., et al., [*The plenoptic function and the elements of early vision*], vol. 2, Vision and Modeling Group, Media Laboratory, Massachusetts Institute of ... (1991).

[26] Velten, A., Wu, D., Jarabo, A., Masia, B., Barsi, C., Joshi, C., Lawson, E., Bawendi, M., Gutierrez, D., and Raskar, R., "Femto-photography: capturing and visualizing the propagation of light," *ACM Transactions on Graphics (ToG)* **32**(4), 1–8 (2013).

[27] Heide, F., Hullin, M. B., Gregson, J., and Heidrich, W., "Low-budget transient imaging using photonic mixer devices," *ACM Transactions on Graphics (ToG)* **32**(4), 1–10 (2013).

[28] Kadambi, A., Whyte, R., Bhandari, A., Streeter, L., Barsi, C., Dorrington, A., and Raskar, R., "Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles," *ACM Transactions on Graphics (ToG)* **32**(6), 1–10 (2013).

[29] O'Toole, M., Heide, F., Xiao, L., Hullin, M. B., Heidrich, W., and Kutulakos, K. N., "Temporal frequency probing for 5d transient analysis of global light transport," *ACM Transactions on Graphics (ToG)* **33**(4), 1–11 (2014).

[30] O'Toole, M., Heide, F., Lindell, D. B., Zang, K., Diamond, S., and Wetzstein, G., "Reconstructing transient images from single-photon sensors," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 1539–1547 (2017).

[31] Tanaka, K., Mukaigawa, Y., Funatomi, T., Kubo, H., Matsushita, Y., and Yagi, Y., "Material classification using frequency-and depth-dependent time-of-flight distortion," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 79–88 (2017).

[32] Satat, G., Tancik, M., Gupta, O., Heshmat, B., and Raskar, R., "Object classification through scattering media with deep learning on time resolved measurement," *Optics express* **25**(15), 17466–17479 (2017).

[33] Tanaka, K., Ikeya, N., Takatani, T., Kubo, H., Funatomi, T., Ravi, V., Kadambi, A., and Mukaigawa, Y., "Time-resolved far infrared light transport decomposition for thermal photometric stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).

[34] Bhandari, A., Kadambi, A., Whyte, R., Barsi, C., Feigin, M., Dorrington, A., and Raskar, R., "Resolving multipath interference in time-of-flight imaging via modulation frequency diversity and sparse regularization," *Optics letters* **39**(6), 1705–1708 (2014).

[35] Gupta, M., Nayar, S. K., Hullin, M. B., and Martin, J., "Phasor imaging: A generalization of correlation-based time-of-flight imaging," *ACM Transactions on Graphics (ToG)* **34**(5), 1–18 (2015).

[36] Jayasuriya, S., Pediredla, A., Sivaramakrishnan, S., Molnar, A., and Veeraraghavan, A., "Depth fields: Extending light field techniques to time-of-flight imaging," in [*2015 International Conference on 3D Vision*], 1–9, IEEE (2015).

[37] Satat, G., Heshmat, B., Raviv, D., and Raskar, R., "All photons imaging through volumetric scattering," *Scientific reports* **6**(1), 1–8 (2016).

[38] Velten, A., Willwacher, T., Gupta, O., Veeraraghavan, A., Bawendi, M. G., and Raskar, R., "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nature communications* **3**(1), 1–8 (2012).

[39] Heide, F., Xiao, L., Heidrich, W., and Hullin, M. B., "Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 3222–3229 (2014).

[40] Kadambi, A., Zhao, H., Shi, B., and Raskar, R., "Occluded imaging with time-of-flight sensors," *ACM Transactions on Graphics (ToG)* **35**(2), 1–12 (2016).

[41] O'Toole, M., Lindell, D. B., and Wetzstein, G., "Confocal non-line-of-sight imaging based on the light-cone transform," *Nature* **555**(7696), 338–341 (2018).

[42] Lindell, D. B., Wetzstein, G., and O'Toole, M., "Wave-based non-line-of-sight imaging using fast fk migration," *ACM Transactions on Graphics (TOG)* **38**(4), 1–13 (2019).

[43] Maeda, T., Wang, Y., Raskar, R., and Kadambi, A., "Thermal non-line-of-sight imaging," in [*2019 IEEE International Conference on Computational Photography (ICCP)*], 1–11, IEEE (2019).

[44] Xin, S., Nousias, S., Kutulakos, K. N., Sankaranarayanan, A. C., Narasimhan, S. G., and Gkioulekas, I., "A theory of fermat paths for non-line-of-sight shape reconstruction," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 6800–6809 (2019).

[45] Kadambi, A., Taamazyan, V., Shi, B., and Raskar, R., "Polarized 3d: High-quality depth sensing with polarization cues," in [*Proceedings of the IEEE International Conference on Computer Vision*], 3370–3378 (2015).

[46] Kadambi, A., Taamazyan, V., Shi, B., and Raskar, R., "Depth sensing using geometrically constrained polarization normals," *International Journal of Computer Vision* **125**(1-3), 34–51 (2017).

[47] Kalra, A., Taamazyan, V., Rao, S. K., Venkataraman, K., Raskar, R., and Kadambi, A., "Deep polarization cues for transparent object segmentation," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], (2020).

[48] Schechner, Y. Y., Narasimhan, S. G., and Nayar, S. K., "Instant dehazing of images using polarization," in [*Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*], **1**, I–I, IEEE (2001).

[49] Schechner, Y. Y., Narasimhan, S. G., and Nayar, S. K., "Polarization-based vision through haze," *Applied optics* **42**(3), 511–525 (2003).

[50] Miyazaki, D., Tan, R. T., Hara, K., and Ikeuchi, K., "Polarization-based inverse rendering from a single view," in [*null*], 982, IEEE (2003).

[51] Atkinson, G. A. and Hancock, E. R., "Recovery of surface orientation from diffuse polarization," *IEEE transactions on image processing* **15**(6), 1653–1664 (2006).

[52] Taamazyan, V., Kadambi, A., and Raskar, R., "Shape from mixed polarization," *arXiv preprint arXiv:1605.02066* (2016).

[53] Cui, Z., Gu, J., Shi, B., Tan, P., and Kautz, J., "Polarimetric multi-view stereo," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 1558–1567 (2017).

[54] Ghosh, A., Chen, T., Peers, P., Wilson, C. A., and Debevec, P., "Circularly polarized spherical illumination reflectometry," in [*ACM SIGGRAPH Asia 2010 papers*], 1–12 (2010).

[55] Baek, S.-H., Zeltner, T., Ku, H. J., Hwang, I., Tong, X., Jakob, W., and Kim, M. H., "Image-based acquisition and modeling of polarimetric reflectance," *ACM Transactions on Graphics (Proc. SIGGRAPH 2020)* **39**(4) (2020).

[56] Treibitz, T. and Schechner, Y. Y., "Active polarization descattering," *IEEE transactions on pattern analysis and machine intelligence* **31**(3), 385–399 (2008).

[57] Maeda, T., Kadambi, A., Schechner, Y. Y., and Raskar, R., "Dynamic heterodyne interferometry," in [*2018 IEEE International Conference on Computational Photography (ICCP)*], 1–11, IEEE (2018).

[58] Tanaka, K., Mukaigawa, Y., and Kadambi, A., "Polarized non-line-of-sight imaging,"

[59] Lindell, D. B., Wetzstein, G., and Koltun, V., "Acoustic non-line-of-sight imaging," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], 6780–6789 (2019).

[60] Saragadam, V. and Sankaranarayanan, A. C., "Krism—krylov subspace-based optical computing of hyperspectral images," *ACM Transactions on Graphics (TOG)* **38**(5), 1–14 (2019).

[61] O'Toole, M. and Kutulakos, K. N., "Optical computing for fast light transport analysis.," *ACM Trans. Graph.* **29**(6), 164 (2010).

[62] Hirsch, M., Sivaramakrishnan, S., Jayasuriya, S., Wang, A., Molnar, A., Raskar, R., and Wetzstein, G., "A switchable light field camera architecture with angle sensitive pixels and dictionary-based sparse coding," in [*2014 IEEE International Conference on Computational Photography (ICCP)*], 1–10, IEEE (2014).

[63] Buckler, M., Jayasuriya, S., and Sampson, A., "Reconfiguring the imaging pipeline for computer vision," in [*Proceedings of the IEEE International Conference on Computer Vision*], 975–984 (2017).

[64] Sitzmann, V., Diamond, S., Peng, Y., Dun, X., Boyd, S., Heidrich, W., Heide, F., and Wetzstein, G., "End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging," *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018).

[65] Sun, H., Dalca, A. V., and Bouman, K. L., "Learning a probabilistic strategy for computational imaging sensor selection," *arXiv preprint arXiv:2003.10424* (2020).

[66] Poh, M.-Z., McDuff, D. J., and Picard, R. W., "Non-contact, automated cardiac pulse measurements using video imaging and blind source separation.," *Optics express* **18**(10), 10762–10774 (2010).

[67] Nowara, E. M., Marks, T. K., Mansour, H., and Veeraraghavany, A., "Sparseppg: towards driver monitoring using camera-based vital signs estimation in near-infrared," in [*2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*], 1353–135309, IEEE (2018).

[68] Nowara, E. and McDuff, D., "Combating the impact of video compression on non-contact vital sign measurement using supervised learning," in [*Proceedings of the IEEE International Conference on Computer Vision Workshops*], 0–0 (2019).

[69] Raghu, M. and Schmidt, E., "A survey of deep learning for scientific discovery," *arXiv preprint arXiv:2003.11755* (2020).

[70] Willard, J., Jia, X., Xu, S., Steinbach, M., and Kumar, V., "Integrating physics-based modeling with machine learning: A survey," *arXiv preprint arXiv:2003.04919* (2020).

[71] Maeda, T., Satat, G., Swedish, T., Sinha, L., and Raskar, R., "Recent advances in imaging around corners," *arXiv preprint arXiv:1910.05613* (2019).

[72] Bhandari, A. and Raskar, R., "Signal processing for time-of-flight imaging sensors: An introduction to inverse problems in computational 3-d imaging," *IEEE Signal Processing Magazine* **33**(5), 45–58 (2016).