Intelligent Imaging

A new imaging paradigm at the edge

For decades deep learning technology lay dormant due to limitations in hardware and data availability. Recent advancements in hardware have allowed this technology to become commercially viable, and has led to its subsequent disruption of numerous fields over the past several years. Meanwhile, the number of image sensors we interact with and depend on is growing at a fascinating rate, and so too are our expectations of these sensor's performance and capabilities. At the intersection of these technologies lies an opportunity to unify sensing with computation such that a new imaging paradigm is born - intelligent imaging. Intelligent imaging enables new use-cases, the highest fidelity signals possible, and a new way to view the world.

I. PROBLEM STATEMENT

Cameras are everywhere. In our smartphones, in our vehicles and in our cities and homes. These imaging systems are as diverse as they are ever-evolving, and not one is perfect. Today we can find impressively compact and low-powered systems that push our fabrication technology to its limits. This need for compact form factor typically leads to a tradeoff for pure image fidelity, which exacerbates noise and imperfections found in even the highest quality of image sensors and electronics.

Noise and other imperfections are as complex and unique as the imaging system that produces them. Traditionally, algorithms meant to combat these imperfections were crafted in a "one size fits all" manner and were deployed on computing hardware that was either too large and inefficient, or too rigid and expensive.

II. BACKGROUND

Imaging imperfections

Ever since the digital revolution in imaging technology, engineers have devised solutions to reduce noise and other imperfections with the aim of enhancing the images and videos their systems produced. These solutions existed at various stages such as minimizing noise in the sensor, conditioning the analog signal, or modifying the digital signal. As in any product design, constraints and tradeoffs dictated the appropriate solution to employ, with none being perfect.

> "These imaging systems are as diverse as they are ever-evolving, and not one is perfect."

Computing hardware

As the development of algorithms to address noise in the digitized signal progressed they became increasingly complex. The most complex of these algorithms could - to some degree - reduce photon shot noise, analog read noise, and correct for colour artifacts at low signal to noise (SNR) ratios. However, the cost of this complexity is the additional computational burden.

Engineers could employ physically larger general purpose computing hardware. These CPU-based systems have greedy power requirements, and their modern incarnations have taken on a multitude of other tasks, leading once again to design tradeoffs between these and the ever increasingly power hungry noise reduction algorithms. The physical size of these solutions also meant that they needed to be implemented at an OEM level, rather than at the sensor or module level.

Another option available to engineers was the application specific integrated circuit (ASIC).

These had the advantage of being purpose built for their task and required significantly less power than the general hardware. However, the design and fabrication of these ASICs is costly, and once fabricated cannot be modified or upgraded. Simple features like digital cropping and scaling are readily implemented with such technology at the sensor level, while more advanced functions are typically integrated into a "system on a chip" package that combines several ASIC function blocks into a single package.

Modern software

Deep learning is a relatively new addition to the software engineer's toolbox. Advances in desktop and server grade hardware, coupled with the availability of high-quality data, has led to a resurgence of this machine learning algorithm. The technique has garnered world-wide recognition and hype for its ability to achieve state of the art results in countless areas from object recognition, image classification, natural language processing, and physical modelling.

"The versatility of this one algorithm cannot be understated."

These algorithms are particularly well-suited to imaging tasks, and have recently proved their ability to perform image-to-image tasks, whereby an input image can be transformed into another image for the purposes of enhancement, modification, or creative endeavours. The versatility of this one algorithm cannot be understated. Reducing noise in an imaging system can now be done in a highly specific way, tailored perfectly to the system being designed.

The tradeoff, of course, is that these new methods can require a staggering amount of computation. Typical deep-learning approaches require multiple power-hungry desktop and server class GPUs, and are impractical for embedded use. However, they should not be so easily discounted for edge-use. These algorithms are composed of only a few basic operations, repeated thousands, if not millions, of times, and do not change drastically from one neural network to the next. This makes neural networks a good candidate for specialized hardware. Their amazing abilities and homogenous nature has driven hardware trends over the past several years.

Hardware that is highly specialized to perform neural computation, sometimes called Neural Processing Units (NPUs), have been developed for servers and desktops, for mobile phone SOCs, and are now appearing as standalone chips. This presents a new mobile computing opportunity, one that is low powered and specialized like an ASIC, while also breaking the ASICs limitations by running software which can change the chips behaviour after fabrication, so long as that software is an incarnation of a neural network.

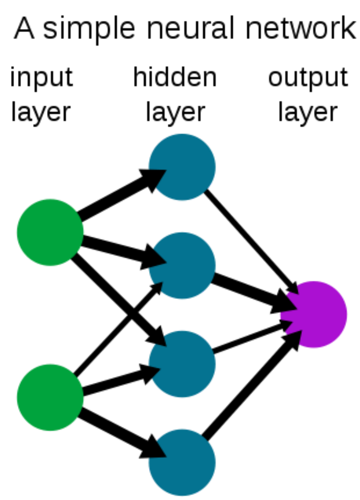

Imaging system designers now have access to a completely new paradigm of software in a compact, low-powered, hardware form factor. What this means is a hardware block that can perform a myriad of tasks, many of which are completely out of reach of even the most complex traditional algorithms.

III. SOLUTION

Enhancing an image, one that suffers from noise and imperfections, can now be performed by a deep learning solution on the edge.

NexOptic's patent-pending Advanced Low Light Imaging Solution (ALLIS™) algorithm *learns exactly how a specific sensor suffers from imperfections, and learns how to optimally correct for these imperfections.* ALLIS provides any imaging system with a state of the art software solution for reducing image noise and imperfections.

To enable ALLIS for an imaging system, a database of sensor information is collected with a representative sensor in a one-time, off-line process we call *profiling*. ALLIS learns from this vast database, seeing millions of example images, and produces a *trained* deep neural network. Once trained, ALLIS is ready for online use.

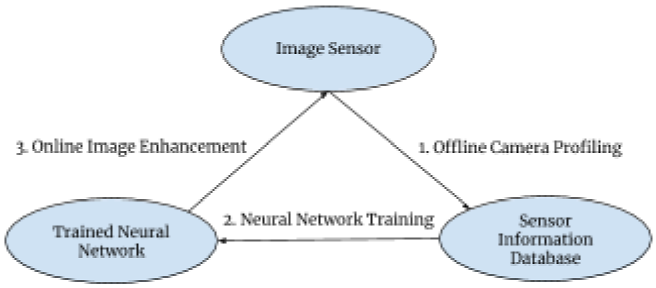

Fundamentally, ALLIS can be deployed onto any hardware platform, but importantly it should leverage today's cutting edge NPU hardware. Matching ALLIS to the right hardware reduces power consumption and increases performance, making ALLIS a powerful edge noise-reduction and image enhancement solution.

ALLIS is a software solution designed to work as a preprocessing step in the imaging pipeline, receiving raw images as input and producing enhanced raw images as output. The exact placement of the solution is dependent on the system designer and specific hardware choice. Modern hardware is offering more flexibility that ever for how and where powerful software executes.

ALLIS can be used as a standalone preprocessor, receiving sensor data directly over industry protocols such as MIPI CSI-2. This offers system designers the ability to integrate ALLIS as a transparent processing block without affecting downstream ISP or processing functionality. In this way ALLIS can be integrated directly into a camera module or system PCB.

Another use case is enabling ALLIS with a coprocessor in a larger system where image processing is performed. Processing in this way is becoming increasingly popular in mobile phones, where blended processing pipelines leverage the speed of dedicated image processors and the flexibility of SoC heterogeneous compute resources. Similar system layouts are found in automotive, IP camera, and medical imaging applications.

IV. CONCLUSION

Deep learning is solving problems across the field of computing science. Imaging in particular is benefiting from this boon of new technological research. Problems that have plagued the imaging community since its conception like image noise, super resolution, object recognition, semantic segmentation and many others are yielding to this seemingly universally applicable algorithm.

These machine learning breakthroughs are driving industry demand for highly specialized hardware accelerators at the edge. Whereas yesterday's digital image sensors allowed machines to see for the first time, today's hardware and software imbue image sensors with vision and perception. This combination of image sensor, hardware and software unlocks new use cases, new sensor operating envelopes, and is driving a trend towards intelligent imaging at the edge.

NexOptic's AI experts bring proven research solutions to industry by training neural networks like ALLIS to perform a wide range of tasks at the edge. ALLIS offers an unrivaled noise reduction and low light enhancement tailored perfectly to a specific image sensor.

AI is expected to transform a number of technology fields and offers product designers a new dimension to add to their offerings. Integrating intelligence will immediately benefit your imaging application and will ensure your products remain relevant in this rapidly evolving field.