CALGARY SCIENTIFIC

Continuous Deep Learning for Visual Systems

James Henry CTO
james.henry@calgaryscientific.com

Overview

Brief overview of deep learning and use cases
Types of AI training (initial, long term collaborative)
Challenges with collaborative learning
Collaborative advanced visualization and continuous deep learning
AI/Human partnerships
Future initiatives

Disclaimer

- I’m not an AI expert
- My apologies in advance for improper use of any terms—please feel free to pass on advice or correct misleading statements… And be kind please. :)
- My interest is in connecting AI to humans and data

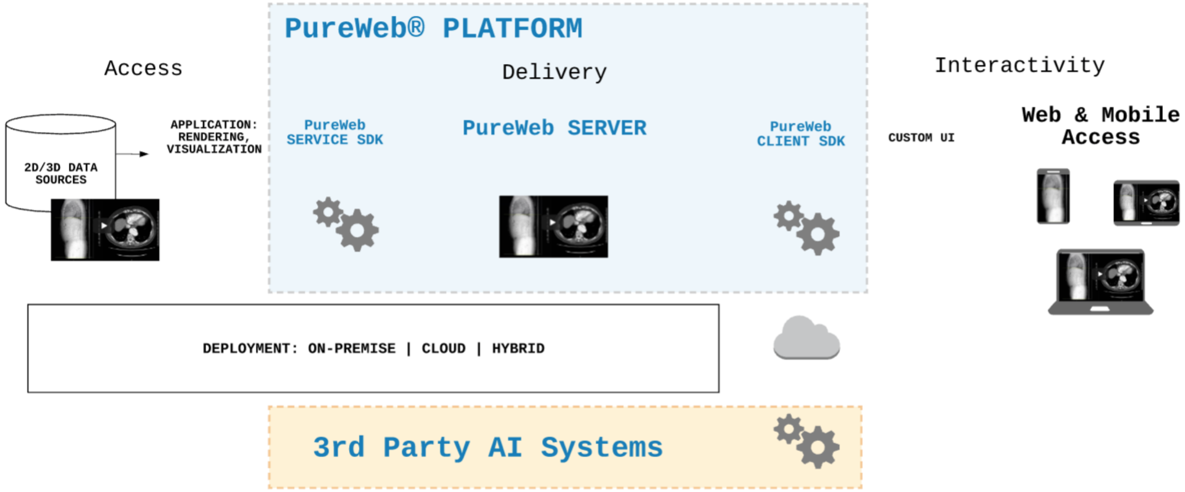

What is Deep Learning (DL)

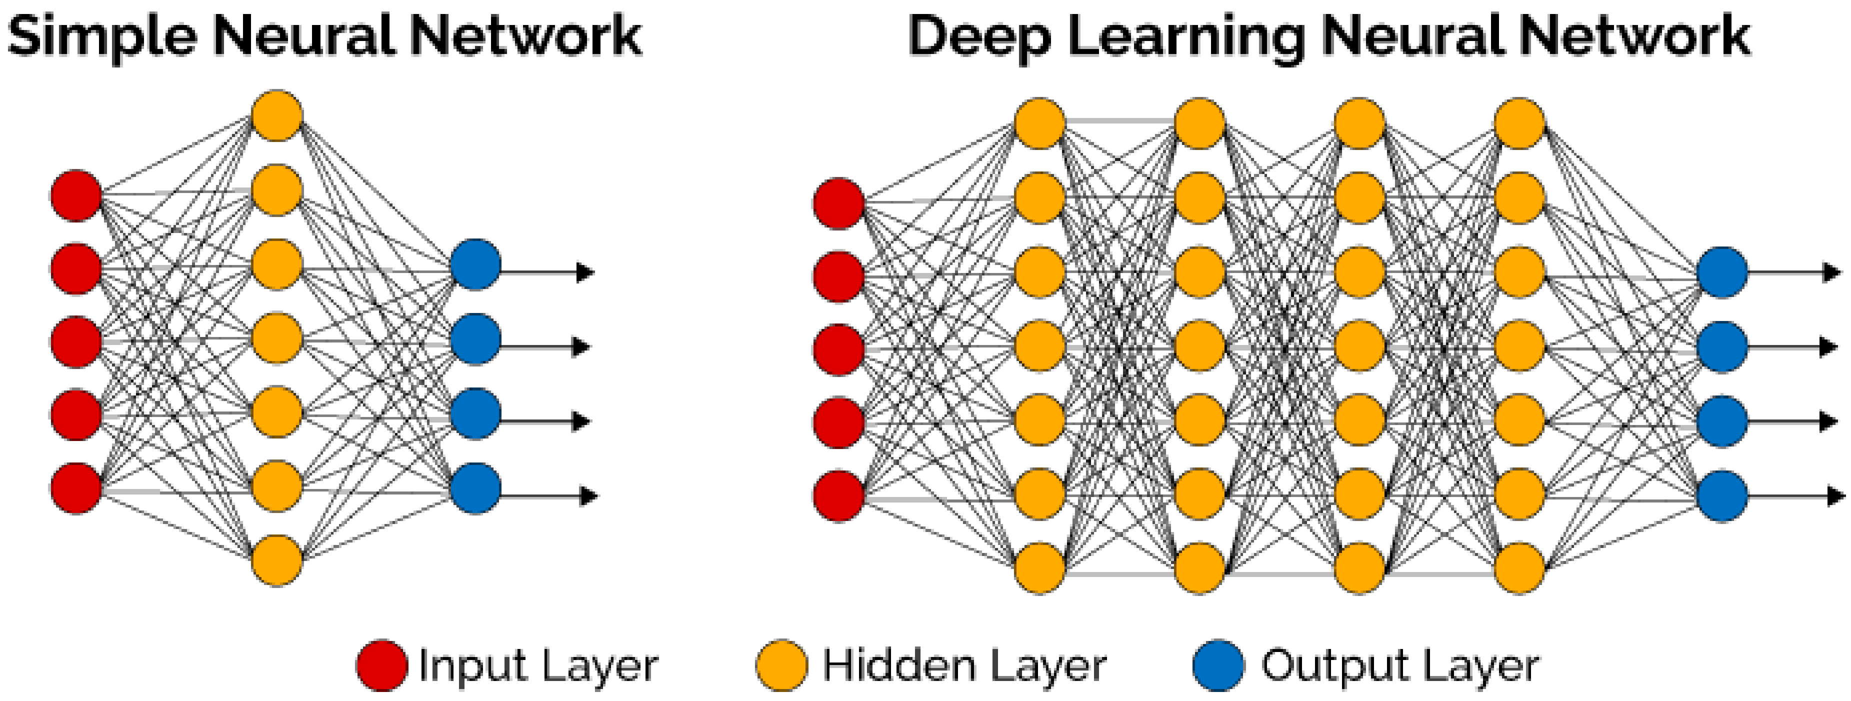

Deep Learning Applications

Some examples

- Language: translation, text to speech, speech to text, handwriting
- Science: protein folding, stellar classification
- Fuzzy problems: sentiment in reviews, multivariate controls

Today's discussion will focus on examples of deep learning algorithms recognizing high value elements of pictures and 3D models

These are **Advanced Visualization** systems (ie "**AV" for short**)

AV System Examples

My deepest experience is with healthcare solutions in radiology and the examples you will see today focus on that area. Real time, mobile, cloud based use cases are particularly familiar to me.

Other industries with similar problems that could benefit from these approaches are:

- Oil and gas (seismic and drilling)
- Design (CAD models, manufacturing use cases)
- Simulation (use of gaming engines for training, product configuration, etc.)

AV Examples

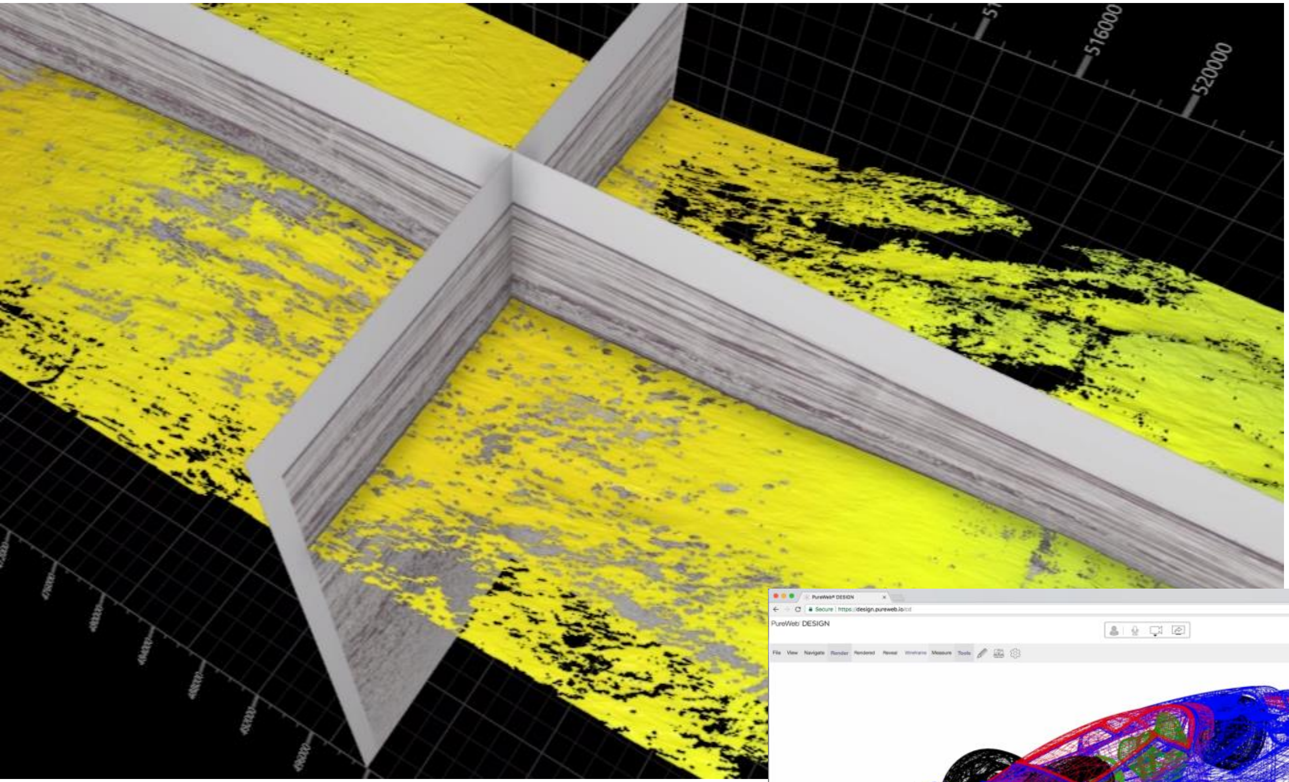
516000
520000

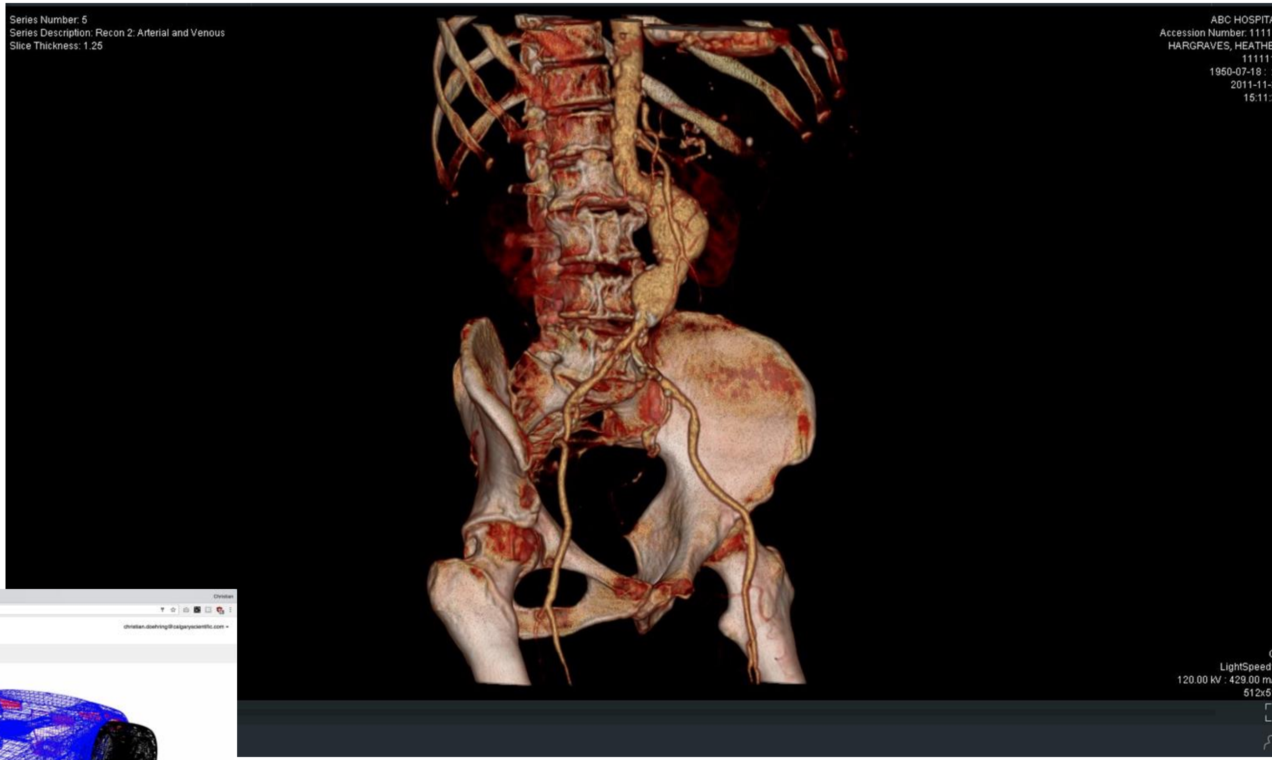
Series Number: 5
Series Description: Recon 2: Arterial and Venous
Slice Thickness: 1.25
LightSpeed

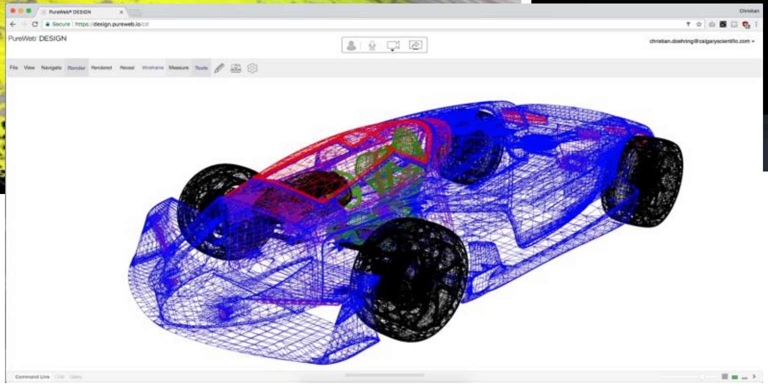
PureWeb DESIGN

Training Stages for AV-based AI Systems

Training the algorithms has distinct phases, for example:

- A data scientist completing initial training
- "On the job training" as the algorithm works in the field

Very different needs for these different modes

Initial Training

Basic process:

- Get data
- Connect it to, say, a DGX station
- Get a data scientist to train and tune it
- Accuracy plateaus eventually

Send your algorithm out to keep learning

Growing the AI system

Think of the algorithm that comes out of the lab as a recent grad; it still needs on-the-job training to hone its abilities

We want to continuously train AI systems to much higher accuracy while they perform useful work

To do that, we need more data, more edge cases and interaction w/ experts to provide results

Interaction with Experts

Problems

Access to experts

Reason for experts to interact

We don't want one expert saying yes/no to 1000s of images

We do want thousands of experts benefiting and incrementally training the AI system

Solution

Get an expert to collaborate with the AI system during a daily workflow to benefit from its automatic processing; then have expert correct bad results

Unfortunately, this means another problem to solve: access

Issues: Combining Experts, AV and AI

Many AV systems are workstation based, which makes collaboration difficult.

Issues

- There are technical challenges with delivering server side rendering results to mobile or remote clients
- On the other hand, putting data on edge devices leads to security breaches, bandwidth limitations, storage and computation challenges
- With sensitive data sets (e.g. medical, military, proprietary designs) it may be necessary to use private, rather than public clouds

Solutions: Combining Experts, AV & DL

Combining experts and AV with continuous deep learning becomes possible if:

- You can centrally locate the data and render complex images on the server side
- Humans can access and interact with the data remotely
- Humans can interact with each other on the data set
- You tie the AI in as a collaborator that provides a valuable service while it learns
- You give a low effort way to flag mistakes the algorithm makes and a way to capture training data for it to learn
- Ideally you can leverage a public cloud to reduce costs; e.g. elastic GPU's

PureWeb® Platform

We developed the PureWeb platform to address the issues with server-side rendering

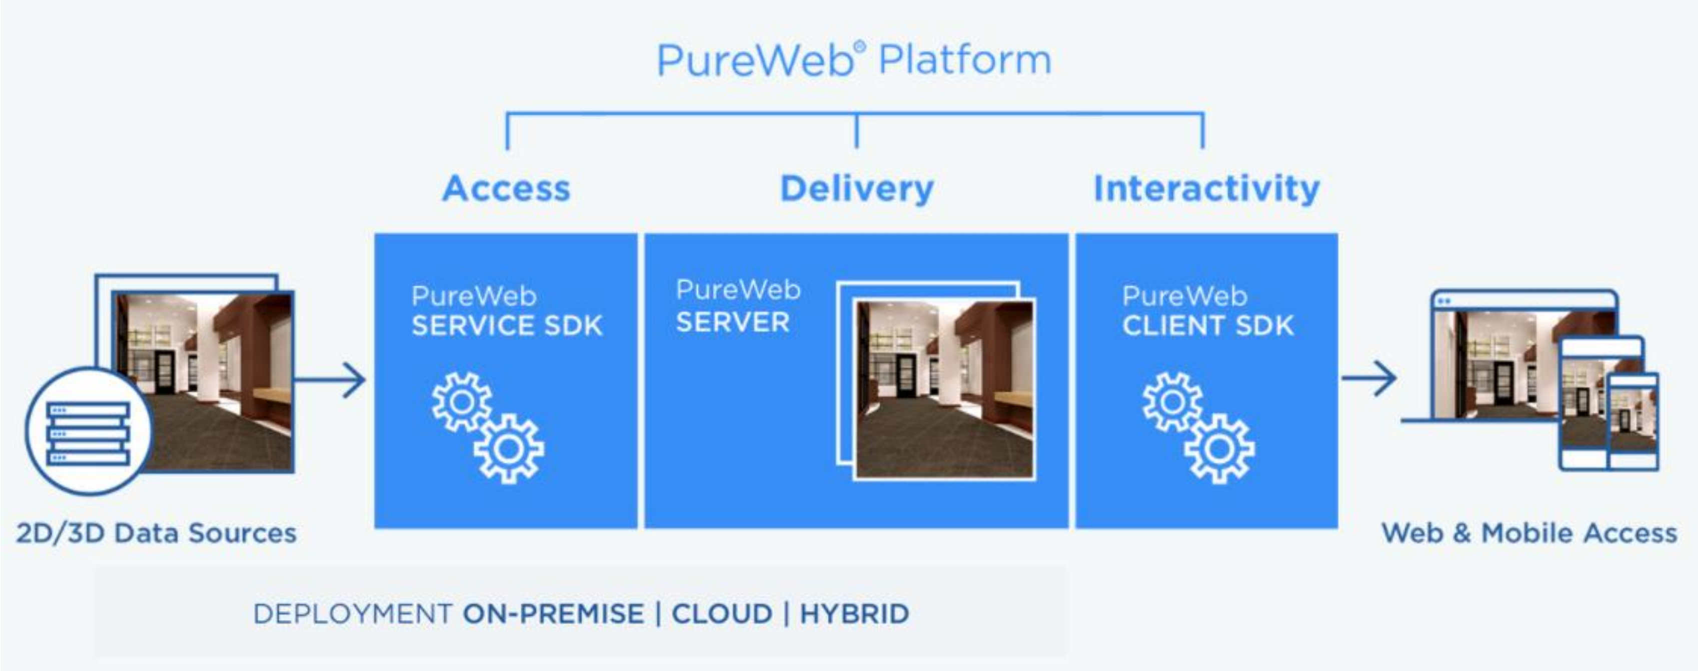

An Example

Humans are great at visual pattern recognition, but we bore easily. We need help on routine work. Some use cases (e.g. radiology) also require very specialized training.

An example problem is screening for diabetic retinopathy. Ophthalmologists routinely capture images of eyes to check for vision irregularities. These images can also show signs of early stage diabetes

If we could turn an AI loose on these images we could potentially intervene to prevent patients from developing diabetes.

Healthcare AV Specific Issues

AI providers don't want to worry about:

- Anonymization
- Site level report customization
- DICOM and HL7 integration
- Deployment
- Finding experts

Our solutions can currently address this in the healthcare industry. It will repeat itself in other industries and need to be solved.

Diabetic Retinopathy

Ophthalmology

- Outline orbit and disk
- Hemorrhage detection

Save time on measuring and outlining, and flag for diabetes

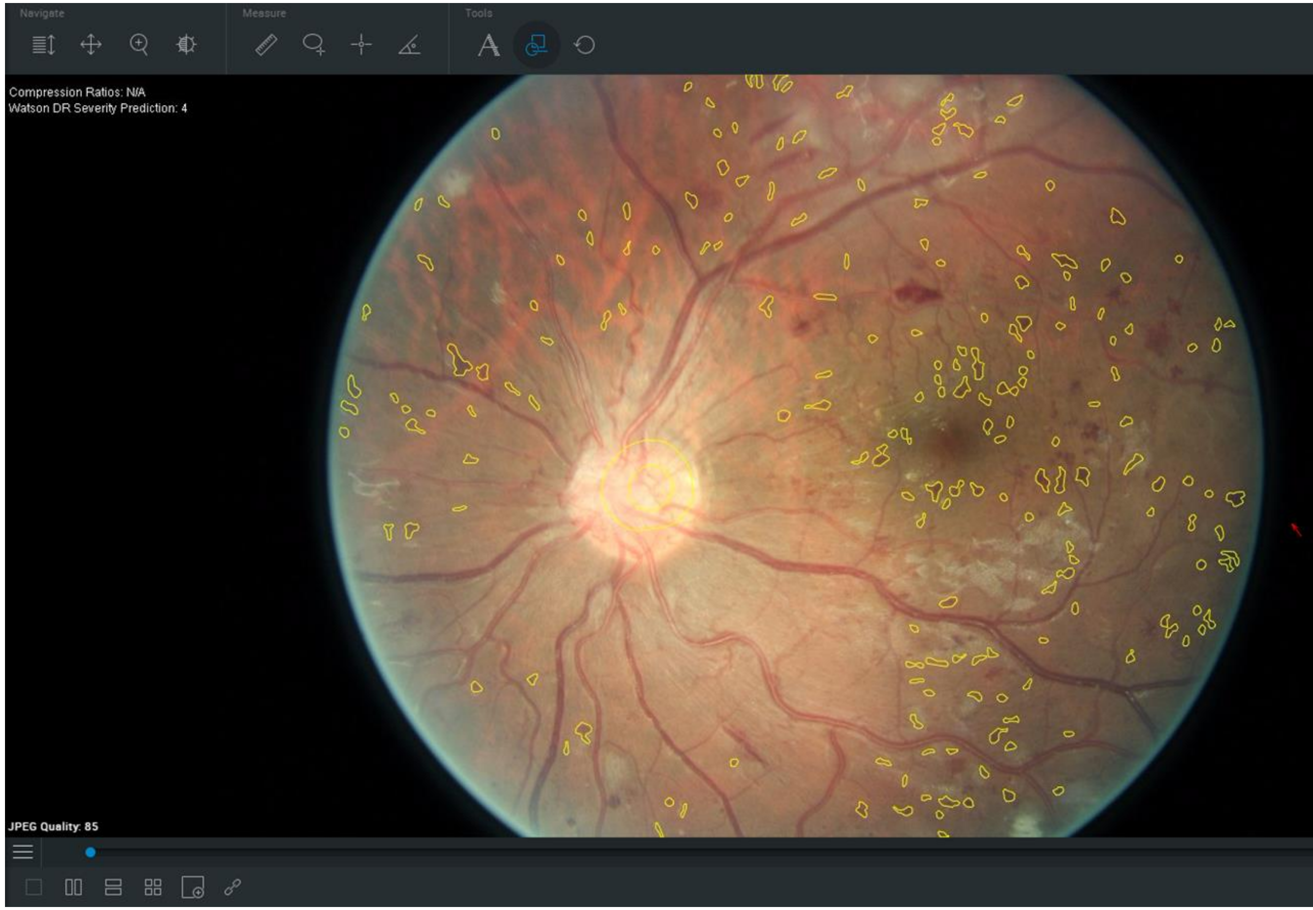

More To Do

Still many pieces to work out in the training use case:

- What degree of feedback is needed? Is "That's wrong" enough?
- What interface elements are required to flag errors? Just a checkbox?
- Is anonymization sufficient to allow sharing of data to train the algorithm?
- How can new data be approved as being appropriate for training?
- Can we pool knowledge from dispersed AI's?
- Storage/bandwidth: can we leave data in place? Can we filter data down?
- This was a 2D example; what are potential 3D examples?
- Legal implications of moving health and financial data to cloud services
- Will we need DL algorithms on private clouds?
- Who will pay?

An Interesting Use Case

A particularly interesting workflow starts to show up as the AI gains in accuracy. What if the AI is correct and the expert is wrong?

Not all "experts" are created equal. This hurdle may actually be pretty low in some cases.

What workflow can we have to check that if the AI is right and the expert is wrong?

At what point does the AI become the gold standard expert and humans have a high bar to correct it?

AI/Human Partnership

In summary, we can see a beneficial partnership in this example:

- Routine pattern recognition is boring and expensive
- AI's are brittle; we need experts in the loop
- Spotting many high reward patterns doesn't happen because of lack of resources/costs (e.g. diabetic retinopathy)
- We can lower costs by automating routine pattern recognition can free up experts to do a higher volume of work and focus on unusual cases
- Lower costs can increase demand if an elastic market

Economist Special Report June, 2017: in Industrial Revolution, number of weavers went up 4x between 1830 and 1900, Luddite concerns notwithstanding

Further Opportunities

There are many more opportunities here:

- Linking multiple AI's to your data set. I've already shown you two connecting. You could get multiple opinions, specialist systems marshalled after an initial screening, or parallel use of AI swarms
- Brokering data and compute connections on clouds to offer SaaS models
- Introducing blockchains to prove provenance and regulate access
- Rendering CAD models of sites against point clouds of actual deployment in digital twin models

CALGARY SCIENTIFIC

Questions?

james.henry@calgaryscientific.com