Pathology AI (Artificial Intelligence) Reference Guide

by Holger Lange and Cris Luengo, Flagship Biosciences Inc.

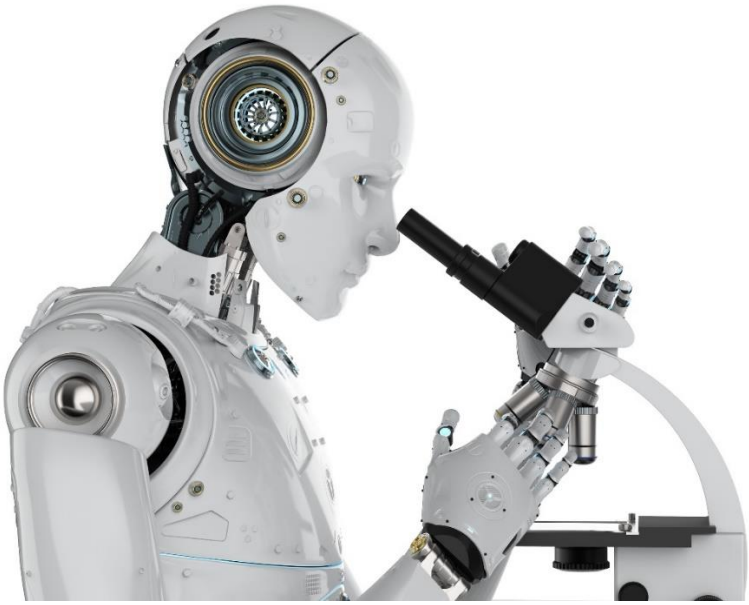

Over the past few years, Deep Learning has created quite a hype about Artificial Intelligence (AI) and Healthcare AI has become a hot topic. We at Flagship Biosciences www.flagshipbio.com have been developing our own Pathology AI system over the last 8 years to solve the most challenging real-world tissue analysis problems across the entire Pharma industry.

You can see a short demonstration of our Pathology AI system for Immuno-Oncology (IO) on YouTube: https://www.linkedin.com/pulse/how-pathology-ai-works-immuno-oncology-short-demo-video-holger-lange/

This reference guide is organized as a collection of independent chapters that each discusses a different key aspect, allowing the reader to select the subject she or he is interested in.

1. Background

We introduce Pathology AI step by step, starting with an introduction to Pathology, going to Digital Pathology and then focusing on machine learning as the key component of a Pathology AI system.

1.1 Pathology

Pathology is the discipline of diagnosing a disease mostly through analysis of tissue, cell and body fluid samples. We focus here on the analysis of tissue. The examination starts with a biopsy. Chemical fixatives, like formalin, are used to preserve the tissue from degradation until the specimen gets to the histology lab. The histology lab processes, embeds, sections, and stains the specimen. Tissue processing removes the water from the tissue and replaces it with a medium, like paraffin wax, that solidifies to create tissue blocks. The tissue blocks are cut into 4 micrometer thick tissue sections using a microtome and are then mounted on a glass microscope slide. Staining is used to give contrast to the otherwise

transparent tissue section and to highlight the features of interest. The most common stain for light microscopy is Hematoxylin and Eosin (H&E) used to see the morphology of the cells. Hematoxylin stains the nuclei blue. Eosin stains the cytoplasm pink. ImmunoHistoChemistry (IHC) uses antibodies to visualize and quantify specific proteins. IHC often uses DAB (diaminobenzidine) to stain the protein in question brown and hematoxylin is used as a counterstain to visualize the nuclei. In light microscopy, the absorption of light by the DAB staining, which is measured as optical density, is propositional to the amount of proteins found in the cells. Appropriate protocols and controls need to be used that control the important variables from biopsy to staining. The histological slides are examined under a microscope by a pathologist. Pathology has several quite distinct subspecialties which correlate to the different tissue types, like prostate, lung or breast. Typically, H&E slides are used to determine the diagnosis, like cancer/no cancer, and a panel of IHC slides are used to determine the right treatment option, requiring an assessment of the protein expressions on the slides in form of a score. There are many different tests in Pathology, all characterized by tissue type and stain, and each of them, while based on complex cell distributions, only reports a single data point per slide, a diagnosis or a score.

1.2 Digital Pathology

Digital Pathology allows to scan the slides and replaces the microscope with a computer monitor. Digital Pathology by itself only provides the convenience of dealing with images instead of glass slides, but by digitizing glass slides to images it enables the use of image analysis and machine learning for tissue analysis.

1.3 Pathology AI (Artificial Intelligence)

A Pathology AI system is a computer program that assists pathologists in their work or provides automated pathology. The key capability of a Pathology AI system is to analyze digital slide images using image analysis and machine learning. Machine learning allows to learn a task from data, like providing a diagnosis or a score, or a subtask, like classifying cells into different cell types. There are many approaches in machine learning, including decision trees, random forests, and deep learning, on which we are going to focus in our discussion.

Over recent years, **deep learning** has created a hype about Artificial Intelligence (AI). Deep learning has overcome major challenges in computer vision, where the feature detection could not be implemented successfully by programming image analysis algorithms. A deep learning network is able to learn highly complex visual features just from the image data, achieving expert human performance. Deep learning requires a lot of data and a lot of processing resources. But recently, with the increase of processing power and in particular the use of GPUs it is now possible to train deep learning networks successfully. AlexNet was the first deep learning network that achieved the major breakthrough in 2012 by significantly outperforming all previous approaches on the ImageNet challenge, a large visual database designed for object recognition software research. Since then every year more efficient and higher performing systems are introduced. As pathology is a visual task it is more than understandable that deep learning is coming to pathology. There has been a Grand Challenge in Biomedical Image Analysis in 2016 and 2017, CAMELYON 16 and CAMELYON 17, on cancer metastasis detection in lymph nodes, which was clearly dominated and won by deep learning. Designing deep learning networks is very challenging, it is no longer about just finding the right hyperparameters, but to design new network topologies, it is an art! Therefore, many applications start out by re-using existing designs that have proven themselves in other applications, like the winner of the CAMELYON 16 challenge re-used the

GoogLeNet that won the ImageNet challenge in 2014. As we know pathology applications are different from a general-purpose image recognition task, a lot could be gained by hand-crafting an appropriate net topology for pathology applications, focusing on cell-data. Only in recent years, academia got started to work on digital slide images, but today most of them are driven by the hype about deep learning. Deep learning has great value where the feature detection presents a challenge for traditional image analysis, but it comes at a price: a) data sets for learning are expensive, b) risk of a bias from the training data and c) no transparency into the decision process.

Decision trees are like hierarchical flowcharts, they are very intuitive for humans and provide transparency into the decision process. Decision trees have many advantages: no data normalization, proper handling of missing data, heterogeneous data (numerical, ordered and categorical), intrinsic feature selection, multi-class, multi-output and fast predictions. Decision trees would be perfect for the use in medical devices, but the question is can they deal with the complexity of analyzing digital slides.

A natural extension of decision trees that provide more complex machine learning approaches are **random forests**. While random forests and decision trees share the same advantages, they give up on the transparency into the decision process, still the important features can be identified. Random forests had a lot of success winning Kaggle competitions that are used to evaluate different machine learning approaches, just before deep learning took the world by storm.

2. A Pathologist-centric Pathology AI system

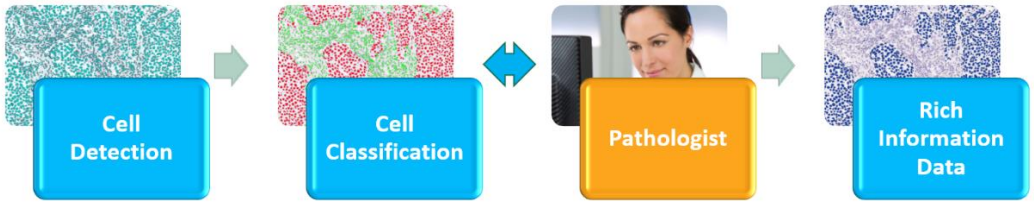

We propose a **system design** of a **pathologist-centric** Pathology AI (Artificial Intelligence) system, illustrated above, where the tissue analysis is semantically segregated into three distinct parts: 1) **cell detection**, 2) **cell classification**, and 3) **measurements which provide rich information data for tissue**, and where the **pathologists provide** their **expertise** in a natural way and the proper **controls** for the system.

We have learned from other applications, where lives are at risk as well, like autonomous driving, that **end-to-end Artificial Intelligence (AI) systems are not the solution**. While astonishing results can be demonstrated with up to 90% performance without having to know anything about the application (quite compelling), it becomes exponentially harder to get to 95% or even the 99% performance where it needs to be.

System designs for applications where lives are at risk are often based on a thorough understanding of the application using **semantic segregation** and **redundancy or control**. Semantic segregation basically means to divide the system into meaningful modules and then to use the right tools for the right tasks. Legal and regulatory considerations favor **assisted vs. automated use** and require **transparency into the decision process**.

Pathologists have been looking at the cells in tissue on histology slides for over 100 years to make their assessments and the characteristics of those cells are representative of the underlying biology. While the implementation details of a Pathology AI system might be too complex to allow transparency, a cell-based representation of the data allows pathologists to understand that data and to be able to interact with the system in a way that is intuitive to them. When we look at markup images of the detected cells in a histology slide, it seems clear that the detected cells contain all the information there is in a histology slide. In this way, **cells** are the perfect **abstraction layer** for a Pathology AI system.

Cell detection can be performed well by **traditional image analysis**, since the imaging process is very controlled and the content in an image is restricted to tissue which has been prepared using a controlled process (unlike other applications where traditional image analysis has its challenges, like autonomous driving, where we look at arbitrary objects in an uncontrolled environment). We are looking for cells that have a certain size, all which have 3 cell compartments (nucleus, cytoplasm and membrane) and which can only be stained by a small number of different stains that have distinct colors. Since computer vision is an area where **deep learning** has achieved major breakthroughs over the last few years, deep learning algorithms can be used in a meaningful way for cell detection.

Cell classification is a critical step that any Pathology AI system needs to accomplish somehow (even end-to-end AI systems, just in a more obscure way) and that ultimately determines the complexity of the machine learning approach and the ability to understand the decision process. We propose a "patient type" based machine learning approach based on cell data, where **pathologists bring in their expertise** by using example regions for the different cell types to train classifiers for different patient types "on the fly". Our approach yields **excellent performance** and provides full **transparency into the decision process**. See the chapter "Machine Learning: Deep Learning vs Decision Trees" for a discussion on our machine learning approach.

Proper controls need to be provided by the **pathologists, verifying** the **cell detection** and **cell classification**!

With the cell detection and cell classification properly verified, all the **measurements that provide the rich information data for tissue** can then just be viewed as simple computer algorithms that can be verified properly during **software verification**.

Providing **rich information data for tissue** as the common data to develop specific **Diagnostics (Dx), Prognostics (Px) and Companion Diagnostics (CDx)** is key to enable Healthcare Big Data for pathology. See the chapter "Healthcare Big Data for Pathology" on what it takes to enable Healthcare Big Data for pathology and what it would look like.

3. Machine Learning: Deep Learning vs Decision Trees

Machine learning, which provides the ability to **learn a task from data** (without the need of being programmed explicitly), is a key component of any Pathology AI (Artificial Intelligence) system. There are many different approaches in machine learning, reaching from simple decision trees to complex deep learning, each with its advantages and disadvantages.

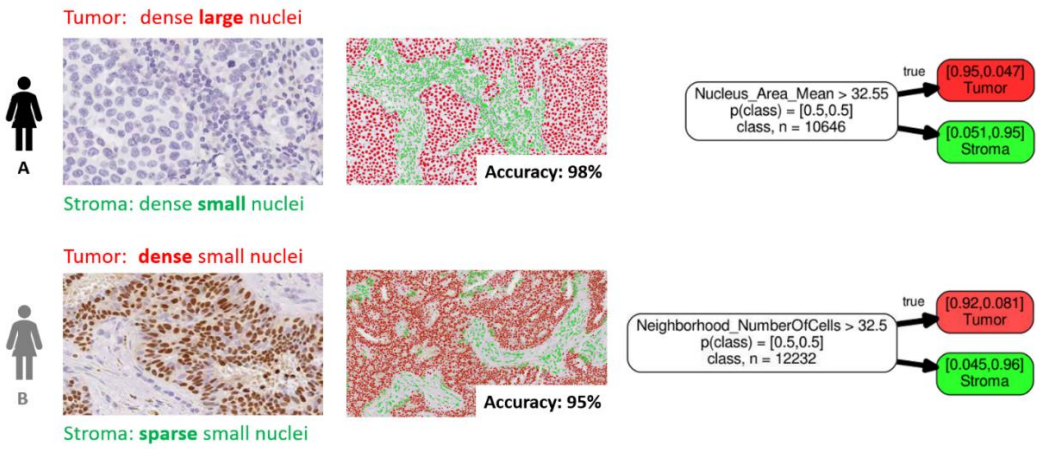

Deep learning, which allows to **learn highly complex visual features**, has created a hype about Artificial Intelligence (AI), as it was able to solve complex computer vision problems that we believed out-of-reach just a few years ago. As pathology is a visual task it is understandable that academia and "pure" technology companies are now working heavily on deep learning approaches for pathology.

The **key problem** for any Pathology AI system are the **variations between different patient types**. In a disease state, no two patient samples look identical. To distinguish between different cell types, which any machine learning system has to accomplish somehow (even if it is hidden in some obscure features in a deep learning network), we notice that the same cell type has different characteristics in different patients, which are often contradictory.

The example above shows two patients that are representative for two different patient types. Patient A has tumor cells that can be described as dense large nuclei and stroma cells that can be described as dense small nuclei. Patient B has tumor cells that can be described as dense small nuclei and stroma cells that can be described as sparse small nuclei. This illustrates nicely what we mean by contradictory characteristics. The stroma cells in patient A have the same characteristics as the tumor cells in patient B: dense small nuclei!

Let's first assume that our machine learning system consists of a **single classifier** that is **pre-trained** using a training set of **histology slide images**.

If we were to create a machine learning system that would be trained only on patients that have the same cell characteristics as patient A, the system would fail when it encounters a new patient type, like patient B. This illustrates the **bias** of machine learning systems that originates from the data used for training. We would need a **lot of training data** to make sure that the machine learning system would be able to learn the characteristics of all patient types properly. Getting a machine learning system from 90% performance to 95% or even 99% performance becomes exponentially harder as remaining exceptional cases are typically hard to come by.

Now if we were to create a machine learning system that would be trained on patients belonging to different patient types, like patient A and patient B, the system would have to learn somewhat **"contradictory" data**. We would need to use a **complex machine learning approach** that would be able to learn highly complex visual features with different contexts.

Obviously, deep learning would be right tool for that job.

Unfortunately, deep learning provides **no transparency into the decision process**, which eventually will have to face some **legal and regulatory hurdles**, as pathology is used to make medical decisions which put human lives at risk.

If the variations between different patient types is the key problem, let's take that out of the equation.

Let's go with a machine learning system which uses **multiple "patient type"-specific classifiers** that **pathologists can train "on the fly"**. See the chapter "Pathologist vs Artificial Intelligence (AI): Competition or collaboration" for a discussion on why you would like to have a pathologist as part of your machine learning system. The training "on the fly" described here is how **pathologists** can best **provide their expertise** in a Pathology AI system.

In pathology there is no critical need to use machine learning to learn the visual features in histology slide. We are not looking at arbitrary objects in an uncontrolled environment, we are looking for cells that have a certain size and that have 3 cell compartments (nucleus, cytoplasm and membrane) and that can only be stained by a small number of different stains that have distinct colors. **Traditional image analysis** can do a good job **detecting cells** and **measuring** a wealth of **biology motivated features** that provide all the information there is in a histology slide.

Let's go with a machine learning system which is **based on cell data**, not pixel data.

That kind of a machine learning system requires **no training data**, yields **excellent performance** and provides **transparency into the decision process**!

Here is how it works.

1) When we encounter the first patient, like patient A, or a patient that belongs to a new patient type (identified in 2), like patient B, we **create a new "patient type"-specific classifier**. A pathologist, using his expertise, identifies a few **example regions** for the different cell types (e.g. tumor and stroma) and **trains a new classifier "on the fly"** that then is used to **classify all the cells on the whole slide**. Proper controls are implemented by having the **pathologist verify the proper classification of the cells**. New example regions are added and the classifier is retrained until the pathologist is satisfied with the cell classification.

2) With any new patient, we first **select the best classifier** from all existing "patient type"-specific classifiers that then is used to classify all the cells on the whole slide. A very simple and robust method that nicely illustrate the selection of the best classifier is to have a pathologist just identify an **example region** for one (or more) cell type(s) and select the classifier that provides the **best classification performance** on those regions. When the pathologist now verifies the proper classification of the cells, he may decide that the classification is not good enough, which means that the new patient belongs to a new patient type and a new "patient-type" specific classifier needs to be created (go to 1).

The example above illustrates nicely that if we were to create different classifiers for different patient types, represented by patient A and patient B, that a simple decision tree using just a single feature with a single threshold would provide excellent performance and easy interpretable decisions. The results obtained by machine learning match nicely what we have seen by eye, the separation between tumor and stroma cells in patient A is based on nuclei size and in patient B on density of cells.

Limiting the machine learning to a specific patient type and using cell data simplifies the machine learning problem considerably, providing excellent performance with **simple machine learning approaches** like **decision trees**, which consist of **easy to understand hierarchical flowcharts**, and only requiring data from very few regions for training. The training is ultra-fast and can be done "on the fly" in an interactive and iterative workflow. A decision tree based on biology motivated features provides **easy interpretable data** to biologists and pathologists and a **meaningful grouping of patients** by patient type.

4. Healthcare Big Data for Pathology

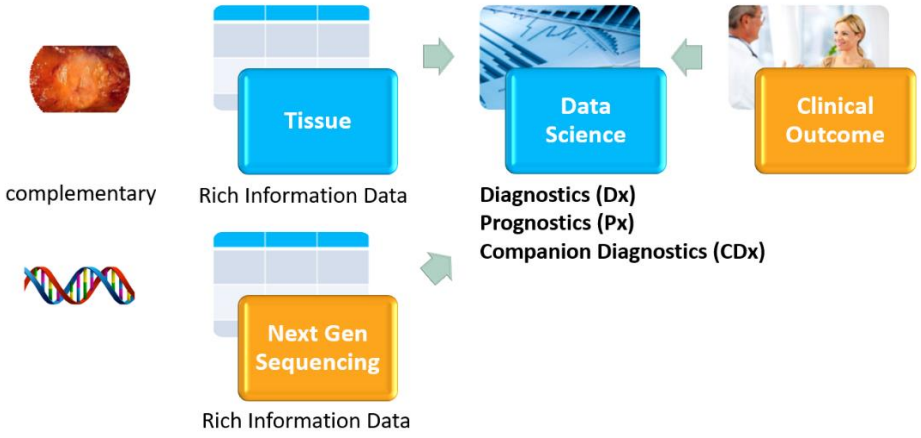

Enabling Healthcare Big Data for Pathology requires moving from single-purpose, single-value pathology data to **general-purpose rich information data for tissue**.

Manual Microscopy, the current standard of care in Pathology where a pathologist examines a histology slide under a microscope, provides only a **single-purpose** (specific indication) **single diagnosis** (e.g. cancer/no cancer) or a **single score** (e.g. 0, 1+, 2+ or 3+) for an entire histology slide. This is a tremendous reduction in data. A histology slide typically has somewhere between 500,000 to over 1,000,000 cells with complex contextual biology (for example: heterogeneity on a slide) and cellular patterns (for example: tissue context as it relates to Immuno-Oncology) that contain significantly more information.

Pathology AI (Artificial Intelligence) systems need to be designed to provide **general-purpose rich information data** for tissue on which any number of single-purpose interpretations or scoring schemes can be implemented. Rich information data for tissue consists of cell-type (e.g. tumor, stroma) specific cell distributions with an abundance of measured features. The key to the value of this rich data is the list of biology based features (e.g. nucleus area size, cell density), which still allow for simple interpretation but capture all relevant information on the slide.

What would **Healthcare Big Data for Pathology** look like?

Imagine a **single test in a clinical lab** (for any given tissue type) where a standard panel of multiple markers provides consolidated rich information data for the tissue. This general-purpose rich information data would support **several specific Diagnostics (Dx), Prognostics (Px) and Companion Diagnostics (CDx)** by simply correlating the rich data to clinical outcome, without requiring creating or running a new test.

Treatment decisions, including the full spectrum of all available drugs (and future drugs, see below), could be **based on such a single test** (solving a major problem in Immuno-Oncology). New **Diagnostic (Dx), Prognostic (Px)** and **Companion Diagnostic (CDx)** capabilities could be **created by clinicians in the field** by correlating existing or emerging health conditions with this "live" clinical database.

For **drug development**, rich information data for tissues would allow better characterization of patient populations, thus creating **smarter tissue and patient selection** strategies which would accelerate and lower the costs of new drug developments.

Rich information data for tissue may also be the **key for considerably simplified regulatory pathways**. In the traditional setting, every tissue type – stain – clinical outcome combination requires a different FDA approval, a costly and time-consuming endeavor. Basically, every Diagnostic (Dx), Prognostic (Px) and Companion Diagnostic (CDx) test needs its own complex FDA approval. However, it may be possible that a **simple FDA clearance per tissue type** would be sufficient for **rich information data** as we are dealing with **general measurements** which are **unrelated to clinical outcome**. During drug development, or when using big data in a clinical setting, you would just need to develop a **scoring scheme that links the already FDA cleared rich information data to clinical outcome**. This would be a much faster and less complex approach.

In the big data realm, traditional single-purpose, single-value pathology data has very limited utility for big data applications. However, general-purpose rich information data allows the full potential of tissue data to be realized in a similar manner as genomics data. General-purpose rich information data for tissue is easy to understand (by a biologist) and easy to access (it's just a database) using any **data science** platform, making it fairly strait forward to combine tissue data with other data sources such as **Next Generation Sequencing (NGS)**, another rich information data source complementary to tissue, and **radiology**.

5. Pathologist vs Artificial Intelligence: Competition or collaboration?

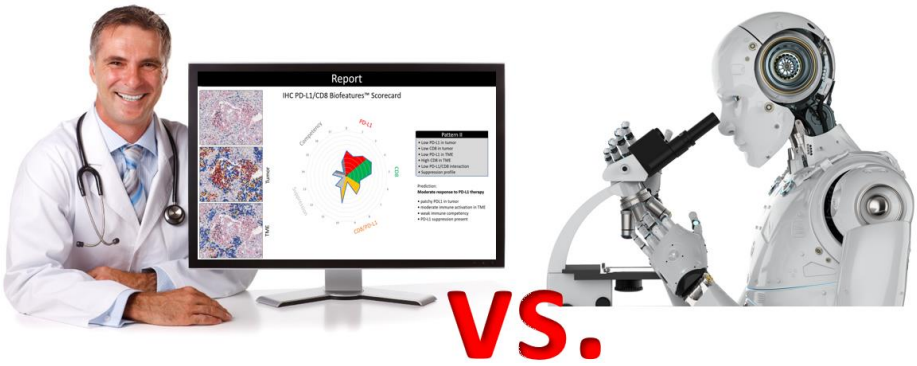

People may think that an **automated Artificial Intelligence (AI) system for pathology that achieves a performance of 90% outperforms pathologists**. But those **pathologists are using microscopes** and the computational part humans are not good at would greatly benefit from the use of computers. So, is that the right comparison?

How good are pathologists actually? Anecdotally, it is interesting to see how pathologists, like nobody else, can agree with other pathologists, even though their diagnosis or scoring was different.

Pathologists agree on what they see in the tissue, based their expertise, but have differences in how to get to the overall diagnosis or score, which is primarily a computational task.

In **Manual Microscopy**, the standard of care and current practice in Pathology, a pathologist "takes a look" at a histology slide under a **microscope**. To start with, microscopes are not FDA cleared medical devices (pre-date the FDA), different microscopes have different optics and even light sources (e.g. blue vs. orange light), and the microscopes in use are often not calibrated properly. Now we are asking the pathologists to assess **500,000 to over 1,000,000 cells** that can have considerable **heterogeneity** across the slide, and reduce that information to a very simple, diagnosis or summary score. For example, in the case of a very simple IHC scoring, we ask the pathologist to determine the **percentage of cells** (to be evaluated against a threshold, e.g. > 10%) of a certain **cell type** (e.g. tumor cells) that have a **staining** (e.g. DAB, that can be collocated with Hematoxylin) in a certain **cell compartment** (e.g. nucleus) that is above an **absolute threshold**. This is a very challenging computational task that obviously will lead to **high inter- and intra- pathologist variations**.

Furthermore, **Immuno-Oncology (IO)**, one of the major advances in drug development in recent years, requires the pathologist now to deal with **more stains**, look at **tissue context** and apply **more complex scoring schemes**. This is becoming **impossible** for a pathologist, if we just give him a microscope.

Obviously, we ask for a lot of things any human has a hard time doing, but a computer can do without a fault. With the **adoption of Digital Pathology**, which enables the analysis of images of histology sides by a computer, it is time to **replace the microscope with the right tools**!

Rather than going "all in" with an automated Artificial Intelligence (AI) system which replaces the pathologists, the use of **Artificial Intelligence (AI) as an aid to pathologists** allows **pathologists** to provide **high-performance high-complexity tissue analysis.**

The right tool for pathologists is a **Pathology AI** (Artificial Intelligence) system that uses **machine learning** for cell classification, where the **pathologists bring in their expertise about the tissue** (identifying the different cell types and verifying the proper cell classification), and have the **computer provide the computational tasks** of counting the cells, calculating objective measurements and complex scoring.

Automated Artificial Intelligence (AI) systems will have a hard time to increase their performance to 95% or even 99%. But that might be required to get close to the **performance of a pathologist using the right tools**.

The key problem for **automated Artificial Intelligence (AI)** systems for pathology are the **variations between different patient types**. In a disease state, no two patient samples look identical. To distinguish between different cell types, which any AI system has to do somehow (even if it is hidden in some obscure features in a deep learning network), we notice that the same cell type has different characteristics in different patients, which are often contradictory. This increases the complexity of the cell classification problem and requires **a lot of training data** to make sure all patient types are included and that there is no bias from the training data. Getting an automated AI system from 90% performance to 95% or even 99% performance becomes exponentially harder as remaining exceptional cases are hard to come by.

The best way, and in our opinion the only viable way, to create enough training data at a reasonable price point, is if the **training data gets automatically generated as part of the standard clinical workflow**.

The **perfect intermediate step** between manual microscopy and automated Artificial Intelligence (AI) is to provide pathologists with the right tools, in form of the **Pathology AI** system described above, which as part of its normal workflow generates an **abundance of free training data** (Pathologist verified cell classification for the whole slide for all slides) that can be used to train an automated Artificial Intelligence (AI) system. As the **performance increases with more training data**, it becomes obvious when the time is right to commercialize an automated Artificial Intelligence (AI) system.

Let's provide pathologists with the right tools, replace the microscope, not the pathologists!

6. A Global Service Model

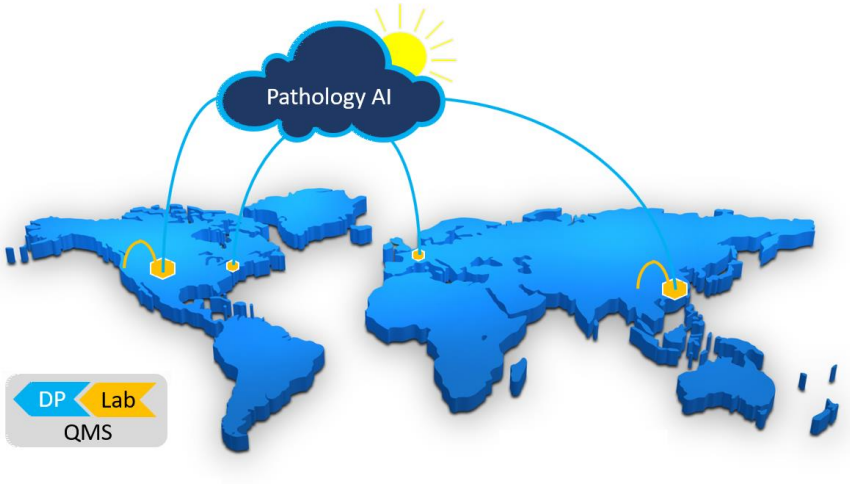

Tissue Image Analysis is a 3-step process: Histology **Lab**, Digital Pathology (**DP**) and **Pathology AI** (Artificial Intelligence).

Histology laboratories create histology slides by processing, embedding, sectioning, and staining tissue specimens taken from biopsies.

Digital Pathology enables a **very simple global service model** for tissue image analysis. Histology slides are digitized to images using slide scanners. The histology slides and the images are created and remain at the local histology labs. **No need to ship tissue or histology slides!** Copies of the images can be shared conveniently and securely around the world and transferred to data centers for tissue image analysis. **No need to create data centers at local lab sites** (minimizing costs and burden on local IT)!

Pathology AI provides tissue data by analyzing images of histology slides using image analysis and machine learning. Since histology slide images are very large and processing can be quite complex, tissue image analysis needs to be performed at **High Performance Computing (HPC) data centers**. A central single HPC data center on-premises (most cost effective) can support a global tissue image analysis service provider from an operational perspective. However now that we are in the days of virtualization and cloud

computing, data centers can easily scale to distributed data centers and data centers in the cloud if needed.

The Pathology AI system itself needs to provide a convenient Digital Pathology interface for transferring images, running tissue image analysis, and viewing and managing the data. Histology labs only need a slide scanner to provide images of the histology slides.

A key challenge for this model is the **regulatory environment**, which depends on both the **country** (e.g. US, Europe, China) and the **intended use** of the tissue data (e.g. **research**, **Clinical Trials**, or clinical **Diagnostics (Dx), Prognostics (Px) or Companion Diagnostics (CDx)**).

Tissue image analysis service providers need to have a thorough understanding of **how to run a histology lab, Digital Pathology and Pathology AI in a regulated environment** (e.g. CLIA, CAP, FDA single-site medical device in the US). They must have an established **Quality Management System (QMS)** with the appropriate **Standard Operating Procedures (SOP)**.

Tissue image analysis service providers can scale from a **single-site central lab model** to a **multi-sites model** by providing any histology lab the **appropriate SOPs** and training to achieve a seamless integration between each site's QMS.

Global service providers need to be able to **operate internationally**. For service providers already operating successfully in the US, which has high regulatory standards, scaling internationally is often straight forward. The service provider will need to work with regulatory experts in the other countries (many consultant companies are available) to comply with the **local regulations** and provide a **localization of their QMS**.

Digital Pathology is a **game changer** for how we think about a service provider **setting up their global network of histology labs**. With the adoption of Digital Pathology, virtually any histology lab in the world with a slide scanner can become part of a global tissue image analysis service provider's network. It comes down to just establishing the global service provider's SOPs. **Sites around the world can quickly be added to the network**, **driven by business opportunities and strategic partnerships.** There is no need to set up and operate a large global network of pre-selected sites. If you want to work with a specific histology lab, simply add it to the network.

7. Pathology AI under a Central Lab Model

A **Pathology AI system** with an intended use for clinical **Diagnostics (Dx), Prognostics (Px) or Companion Diagnostics (CDx)** can be commercialized under a **central lab model** either as a **Lab Developed Test (LDT)** or a **single-site medical device**.

Compared to the development of a commercially-distributed medical device, a **central lab model** can provide **agile** and **cost-effective alternatives** and a **smart intermediate step** (or two) **to a commercially-distributed medical device**. See the chapter "Pathology AI as a Medical Device" for a discussion on what it takes to commercialize a Pathology AI system as a medical device.

Pathologist Cockpit

In a **central lab model**, local clinical laboratories across a country send their tissue samples for testing to a central national laboratory and get the test reports back. **A single-site laboratory can provide an effective distribution channel to an entire country!** In the US, a similar reference lab model is already well established for complex IHC testing. See the chapter "Global Service Model for Tissue Image Analysis" for a discussion on what it takes to **scale** from a single-site central lab model to a **global multi-site model**.

Central laboratories (single- and multi-sites) need to establish a Quality Management System (QMS) to perform Digital Pathology and Pathology AI in a regulated environment. Only the central laboratories need to adopt Digital Pathology and buy a slide scanner. **No dependency on the adoption of Digital Pathology for any other laboratory!** The local laboratories just need a computer and an internet connection to be able to view the images of the histology slides and access the data and reports.

Clinical laboratories in the US can provide a Pathology AI system as a **Lab Developed Test (LDT) under CLIA regulations and CAP standards**.

A LDT can be assembled using different lab-specific components that do not need to be medical devices. Any slide scanner that can be found in the laboratory, and 3rd party software, like Photoshop, can be used to build a validated LDT. **No need to manufacture the Pathology AI system as a medical device!**

Pathology AI systems can move seamlessly from research to a clinical laboratory. Pathology AI applications can be developed as research prototypes until they show clinical utility in an exploratory setting. The prototype can then be brought into the clinical laboratory as a LDT and be distributed under the central lab model. **Implementing a LDT is fast and cost-effective!** The major effort and costs are in the CLIA or CAP validation of the LDT.

A LDT provides the opportunity to gather **clinical utility** and **safety data** in a real-world, but highly controlled setting, that helps to **build confidence in this new technology** for clinicians, regulators, and payors, alike. The clinical data will show if there are any issues before engaging in a costly commercialization as a medical device, therefore a LDT can be used as a **risk mitigation** for the development of the medical device. Since the LDT and the medical device are based on the same

algorithm producing the same data, the **clinical data obtained by the LDT** could be **used to support the regulatory approval of the medical device**.

Another material incremental step from a LDT to a commercially-distributed medical device is a **FDA approved single-site medical device**.

A FDA approved single-site medical device requires that the clinical laboratory adds **design controls** to the development (think assembly) of the LDT (not to the development of its components) as well as it needs to obtain a **regulatory approval** based on **site-specific clinical studies** that can be more comprehensive than a CLIA or CAP validation for the LDT. Still there is no need to manufacture the Pathology AI system as a medical device! The increased rigor in the design controls and validation studies allows to build further confidence in this new technology.

Expanding the use from a FDA approved single-site medical device with proven performance to a FDA approved commercially-distributed medical device, breaks down the regulatory approval into two smaller meaningful steps. Therefore, this is a good **risk mitigation** for the regulatory approval of the commercially-distributed medical device. Assuming that the medical devices share the same intended use, the same medical device definition and especially the same algorithms, the only burden left on the regulatory approval of the **commercially-distributed medical device** should be the manufacturing of the Pathology AI system as a medical device and performing the multi-site reproducibility studies. **A lot of work, time consuming and expensive, but very low risk!** Compare that to developing a medical device from scratch.

When a medical device is needed, but the **business case** for a commercially-distributed medical device is **a challenge**, for example when the patient populations are limited, like in the case of second line therapy indications (now quite frequent in the Immuno-Oncology space), or rare diseases, a commercialization as a single-site medical device under a central lab model can provide a viable alternative.

8. Pathology AI as a Medical Software Device

A **Pathology AI system** with an intended use for clinical **Diagnostics (Dx), Prognostics (Px) or Companion Diagnostics (CDx)** can be commercialized as a **Medical Device**.

Commercialization as a medical device requires a) that the **Pathology AI system** is **manufactured as a medical device**, b) that **clinical studies** are conducted, c) the appropriate **regulatory approvals** are obtained, and d) a **(global) distribution channel** is established.

With the adoption of **Digital Pathology**, clinical laboratories are using slide scanners from different manufacturers. As those slide scanners account for the major investment in Digital Pathology, any Pathology AI system needs to be slide scanner agnostic and be compatible with slide scanners commonly found in clinical laboratories. Nobody wants to buy a new slide scanner just to be able to use a new Pathology AI system. As in radiology, pathology now needs to be able to plug-and-play with different modalities and devices to enable widespread adoption. **Pathology AI systems** need to be their **own medical devices** with a clear and easy path to work with **commodity slide scanners**. A Pathology AI system could be a medical device by itself that delivers reports for pathologists and oncologists, but often clinical laboratories would like to see an integration with their existing **LIS/LIMS**, which most LIS/LIMS today allow to do easily.

A Pathology AI system can be considered **Software as a Medical Device (SAMD)**, using Commercial-Off-The-Shelf (COTS) hardware. To **manufacture medical device software**, a company must have a software engineering team and a QA team with an established **Quality Management System (QMS)** based on **ISO 13485, ISO 14971, IEC 62304** and **FDA QSR**.

Pathology AI **Diagnostics (Dx)** and **Prognostics (Px)** products can be developed by "just" **technology companies**, no IVD manufacturer needs to be involved when common H&E staining or already approved IHC assays are used. The key is to get access to patient outcome data for the development and the clinical trials, which typically implies partnerships with **Academic Medical Centers (AMC)**. The entire burden on the development and commercialization is on the technology company.

Conversely, Pathology AI **Companion Diagnostics (CDx)** products are **driven by Pharma** companies developing therapeutics who need a test for patient selection as part of their drug label, thereby requiring a continuous path from their clinical trials to a Companion Diagnostics (CDx). The Pharma companies typically provide the data for the development and conduct the clinical trials required for the regulatory approvals as part of a therapeutic/diagnostic co-development. However, Companion Diagnostics (CDx) products typically also involve a partnership with **IVD (In vitro Diagnostics) manufacturers** who provide the wet **assays** and have the appropriate **distribution channels**.

Building a Pathology AI system from scratch is quite an endeavor, as a lot of expertise and experience needs to go into its development. But often, as seen in the case of a Companion Diagnostics (CDx), some **existing technology has shown clinical utility in an exploratory setting** and "just" needs to be manufactured as a medical device with a simplified user interface and locked-down algorithms. At this moment, it becomes important to control the software, relying on **in-house software development**, as opposed to using 3rd party software. Using a **validated prototype** (e.g. LDT under CLIA/CAP) as the basis for a medical device development **simplifies and de-risks the development lifecycle** considerably, which at that point can become a simple waterfall model. Being able to manufacture a medical device that provides the same data as its prototype during a clinical trial opens the door for a **product development**

parallel to the clinical trials, such that a simple bridge study can be used to allow the medical device to use the study data from the clinical trial for its regulatory approval.

Once a Pathology AI system has been manufactured as a medical device it provides a platform to which new applications for **Diagnostics (Dx), Prognostics (Px) or Companion Diagnostics (CDx)** can be added very quickly (basically only the algorithms), ultimately building a comprehensive **Pathologist cockpit**.

To run **clinical trials** and get the appropriate **regulatory approvals** for a medical device, a company must have the expertise to conduct clinical trials and to obtain regulatory approvals for IVDs globally, ideally having established **Standard Operating Procedures (SOP)** for **Good Clinical Practices (GCP) for Clinical Trials**.

Technology companies need to build their own **global distribution channels** into the clinical laboratories, or work with partners, like **IVD manufacturers**, that already have the appropriate **distribution channels**.

The **commercialization** of a **medical device** is **very costly**! The US anatomic pathology **market** for a medical device based on the current reimbursement model is relatively small, we estimate it to about **$11 million**, which then again is shared by multiple manufacturers. The problem is that the anatomic pathology market is very segmented by subspecialties, which correspond to different tissue types (e.g. breast), for each of which a list of different tests (e.g. H&E diagnosis, IHC Her2, ER, PR) exists, which typically correspond to different stains. This creates a myriad of "**tissue – stain – clinical outcome"** specific tests that each must become a separate **medical device**. The **business case** to commercialize a Pathology AI system as a medical device is **a challenge**. See the chapter "Pathology AI needs a new business model" on a discussion about the business case for Pathology AI.

Interestingly, all **existing tissue image analysis medical devices for pathology** (Computer assisted IHC Her2, ER, PR, …) have been developed by Digital Pathology manufacturers, but with a completely **different business case** in mind, to be able to market their Digital Pathology equipment into the clinical market. No real Pathology AI medical device exist today.

9. Pathology AI needs a new business model

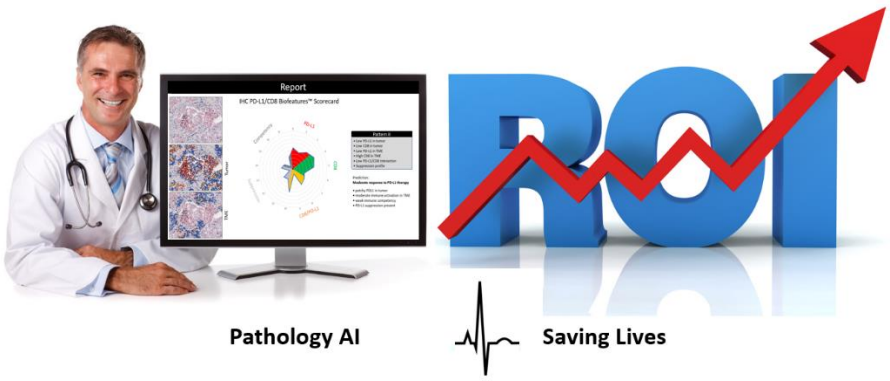

The true barrier for Pathology AI is the business model, not the technology. To get Pathology AI into the clinical laboratories, payers need to provide a **value-based model** that creates a viable business case.

We estimate the **US anatomic pathology market** for **tissue image analysis** based on the current **reimbursement model** (e.g. computer assistance = CPT 88361 – CPT 88360 = $7-8) to be about **$550 million** with about **$7-8 per test**, even though it is unlikely that CMS (Centers for Medicare & Medicaid Services) is going to just add $550 million to their reimbursements.

The problem is that the anatomic pathology market is segmented into subspecialties, which correspond to different tissue types (e.g. breast), each of which has a list of different tests (e.g. H&E diagnosis, IHC Her2, ER, PR), which typically correspond to different stains. This creates a myriad of "**tissue – stain – clinical outcome" specific tests**, each with its own little **market segment** that we estimate to be on average about **$11 million** (e.g. breast – IHC HER2 – score), which is shared by multiple manufacturers.

Pathology AI is **dependent on the adoption of Digital Pathology**, which by itself does not have a tangible business case (unlike radiology). Depending on whether a laboratory has a scanner, any additional reimbursement for computer assistance may need to **fund the purchase of the Digital Pathology equipment as well**. Ultimately Pathology AI will **drive the adoption of Digital Pathology**, providing it with a Return-On-Investment (ROI)!

When you consider the costs associated with building and commercializing a Pathology AI system as a medical device, this **business case** becomes **a challenge**! See the chapter "Pathology AI as a Medical Device" for a discussion on what it takes to commercialize a Pathology AI system as a medical device.

What is the killer app?

Applications that provide the **same results as pathologists using a microscope**, just providing **better consistency** or **saving time**, **make almost no difference in the market**!

We have seen this very clearly with the tissue image analysis IHC HER2 test for breast cancer, the poster child for these kinds of applications. The adoption started very strong between 1998 and 2002 when the additional reimbursement was very high, about an additional $170 per test, and by 2002, about 450 ACIS systems (the first commercial product) were placed. This reimbursement changed to under $60 in 2003, and in 2007, 5 years later, only 250 ACIS systems were still in the market. Today, the additional reimbursement is under $10 per test.

Interestingly, several additional tissue image analysis IHC HER2, ER, PR, etc. medical devices have been developed over the years, all by Digital Pathology manufacturers, who have a completely different business case in mind, to market their Digital Pathology equipment into the clinical market. Given that there was a predicate device, this allowed for an easy route.

Rare event detections, including pathogens, like Acid Fast Bacillus (AFB), or cellular patterns, like mitotic figures, which could save pathologists a lot of time, didn't' even get that far. Why would the diagnosis of cancer, the latest application that everybody is talking about, be any different? After all, the gold standard is a pathologist using a microscope, why change to potentially lose money?

The adoption of Pathology AI, under the current reimbursement model, will only be driven by "**microscope impossible" tests that require a pathologist to use Pathology AI**. Today, **Immuno-Oncology (IO)** is the **killer app** with a massive business case behind it. There is an extensive need for tissue context data that other modalities, like next generation sequencing, cannot provide. The required analysis of the tissue is far too complex for a pathologist, if we just give him a microscope.

What is the right business model? Where is the value?

Pharmacogenomics allow us to identify the patients who are more likely to respond to particular therapies or who require dose modifications. Stratification of clinical trials, even retrospective, boosts efficacy and eliminates toxicity. How much is that worth?

Prescribed cancer treatments are effective in only about 25% of cancer patients, making them inefficient, expensive and detrimental to patient health. Adverse drug reactions annually in the US alone, account for **100,000 patient deaths**, **$100 billion healthcare costs**, and is the **4th leading cause of mortality**. Between 1997-2004 19 **drugs** were **removed from the market** (based on 2008 data).

Continuing with our example of HER2, the drug Herceptin acts on the HER2/neu (erbB2) receptor. In normal cells, HER2 controls aspects of cell growth and division. When activated in cancer cells, HER2 accelerates tumor formation. About 20-30% of breast cancers overexpress HER2 meaning those patients may be candidates for the drug, which costs about $70,000 US dollars for a full course of treatment. One of the significant complications of Herceptin is its effect on the heart, it is associated with cardiac dysfunction in 2-7% of cases. An IHC HER2 test that identifies the 20-30% of patients who would benefit from a $70,000 treatment and eliminates the risk of a cardiac dysfunction in the 2-7% of the 70-80% of patients that would not benefit from the treatment would just ***save the healthcare system*** *about* ***$50,000 per test****, not to mention the saved lives.*

The true opportunity for Pathology AI is **personalized medicine** with **big data**. This means that we could run a **single test in a clinical laboratory** (for any given tissue type), a standardized panel with multiple markers that provides **rich information data for tissue**. **Treatment decisions** that include the **full spectrum** of all available and future **drugs** could be based on that **single test**. New **Diagnostics (Dx), Prognostics (Px)** and **Companion Diagnostics (CDx)** could be **created by clinicians in the field**, correlating existing or emerging health conditions with this clinical "live" database. **Drug developments** could be done **faster** and **cheaper** with a smarter patient selection that could be achieved through better characterization of a patient population using the rich information data from that test. **Diagnostics (Dx), Prognostics (Px)** and **Companion Diagnostics (CDx)** that were based on that same test could significantly **simplify the regulatory pathways**. See the chapter "Healthcare Big Data for Pathology" for a more detailed discussion on our vision of Healthcare Big Data for Pathology and what it takes to get there.

Pharma companies who want to bring their drugs to market and **payers** who want to improve patient care and lower healthcare costs need to provide a **viable business model that is based on the value** that is provided by **pharmacogenomics** and **big data**.