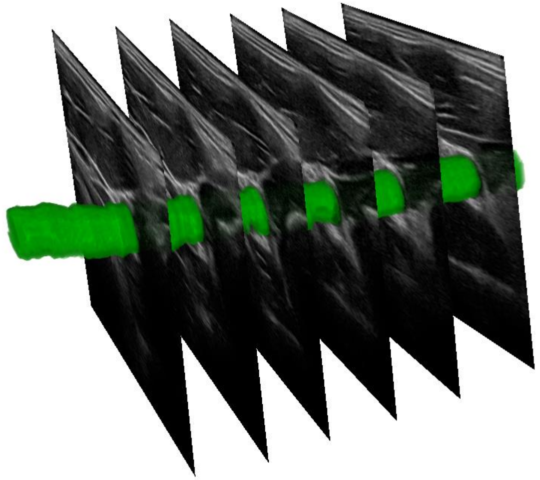

Shaping the Future of Medical Ultrasound Imaging with AI and GPU Computing

GTC 2019 Conference – Session S8712

Raphael Prevost

Senior Research Scientist @ ImFusion

ImFusion

In this session, you will see how AI enables to…

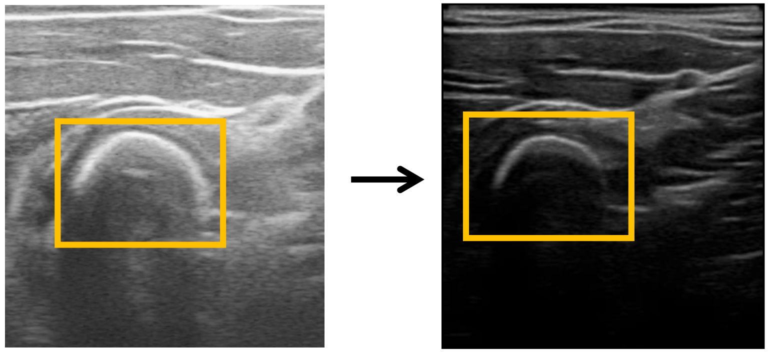

Develop an
auto-focus system…

… transform a video clip
into a 3D volume…

…improve orthopedic
and brain surgery…

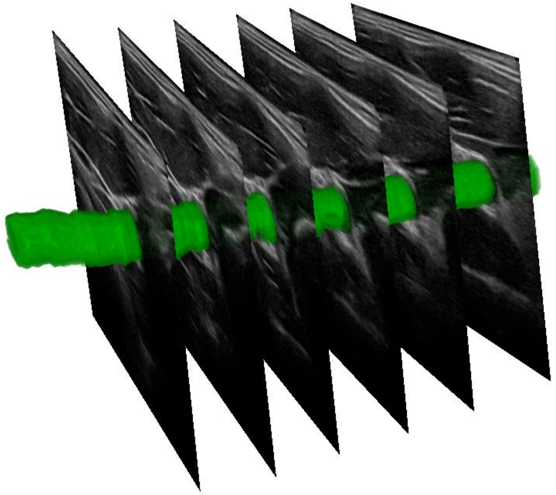

…and reconstruct
my carotid in real-time !

But first, who are we?

ImFusion

- Company founded in 2012
  in Munich, Germany
- Private and independent R&D lab
  in medical imaging and computer vision
- Software framework deployed in various clinical products and used by large companies, start-ups and research labs

What we do

Project Consulting

From feasibility studies to implementation

Research & Development

Solving challenging problems with state-of-the-art algorithms

Software Development Kit

ImFusion SDK serves as an ideal platform for R&D

Implementation & Integration – OEM

Running our software within your medical product

Technology Portfolio

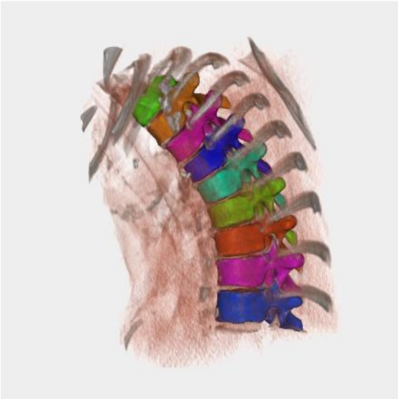
Image Segmentation

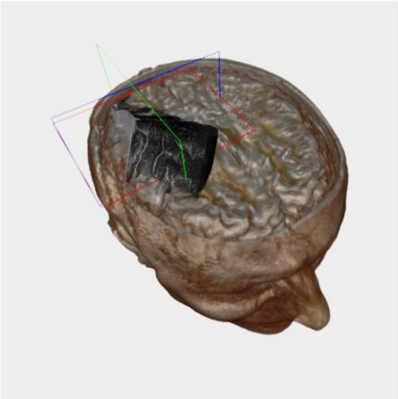
Image Registration

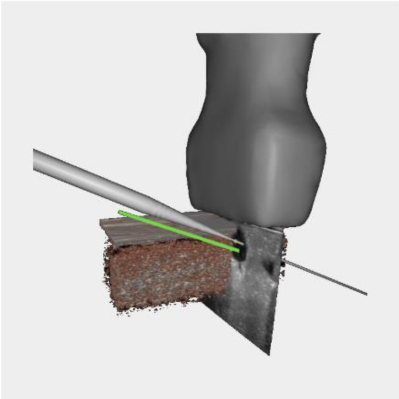
Navigation

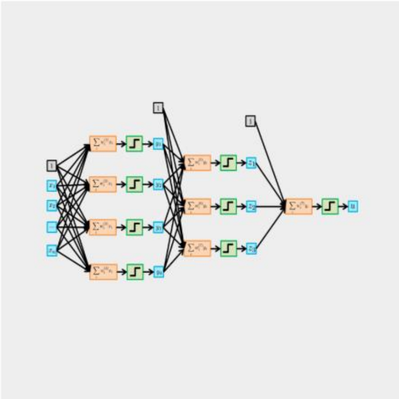
Machine Learning

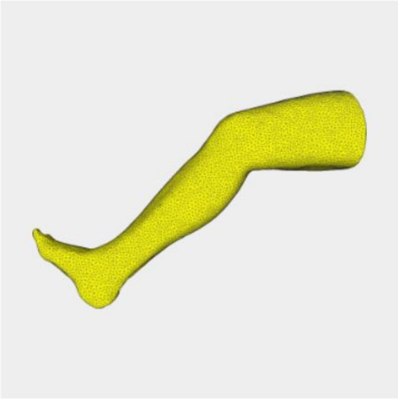
3D Scanning & RGB-D

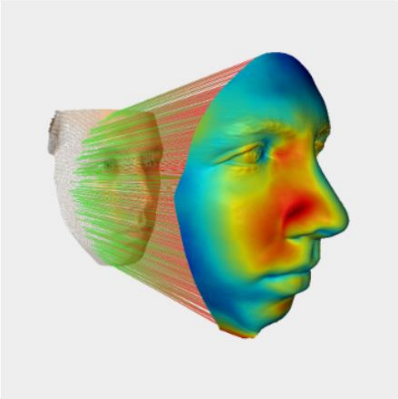
Point Clouds and Meshes

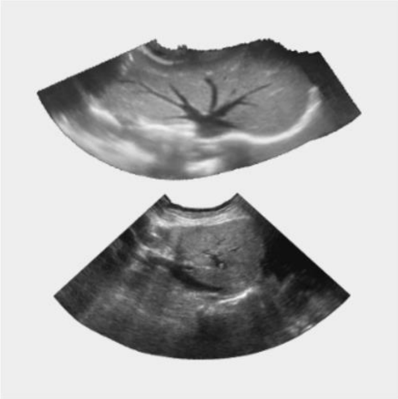
3D Freehand Ultrasound


Cone-Beam CT


Diffusion Tensor Imaging

More...

Part 1
Quick Intro To Ultrasound

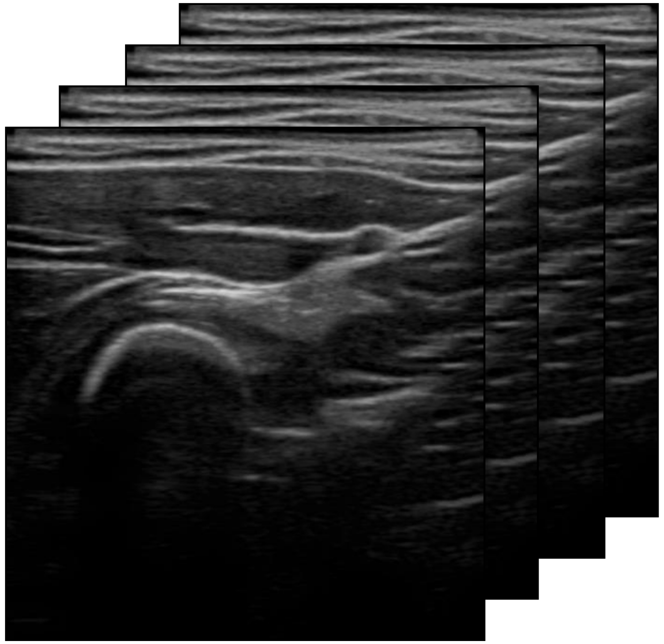

Ultrasound for Medical Applications

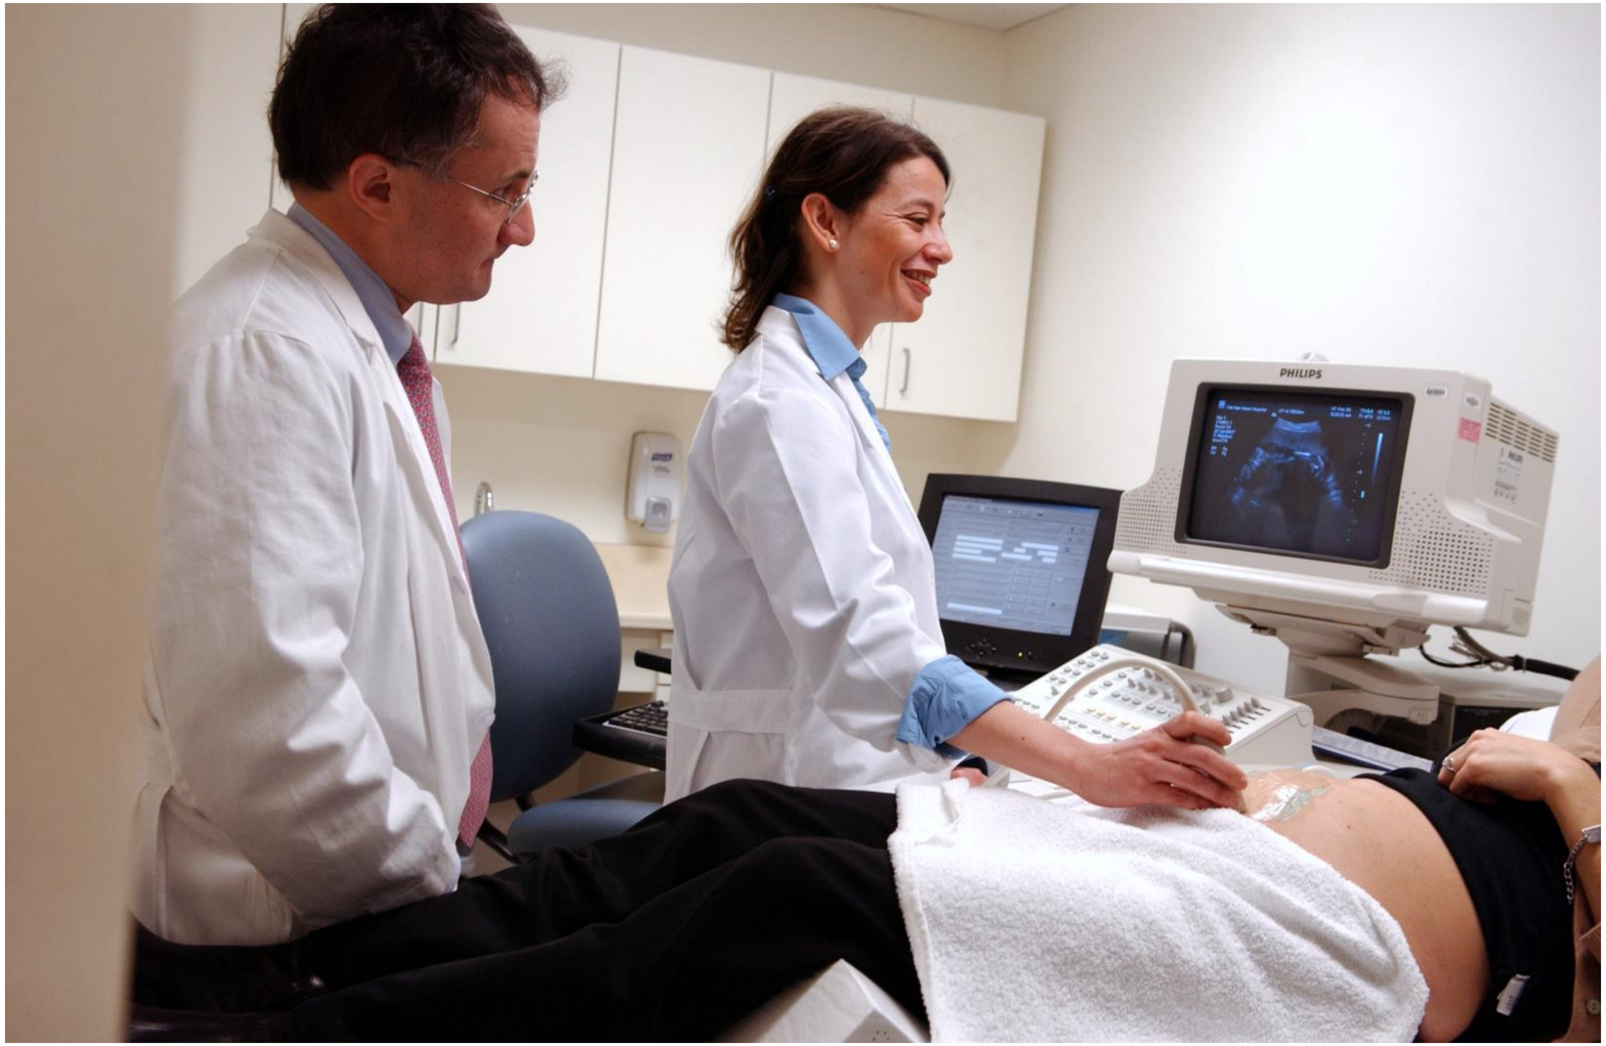

Credit: Yale University

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

Problem #1: US is difficult to acquire

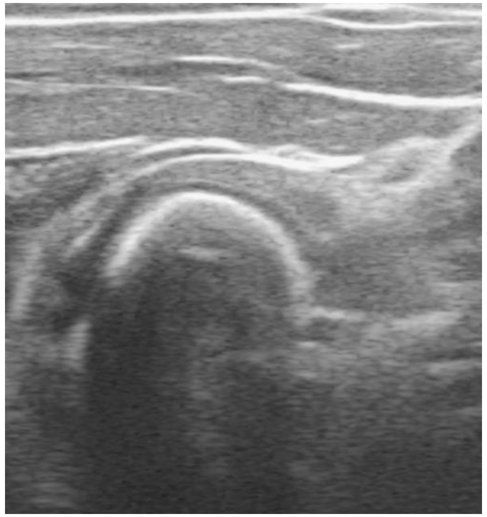

Problem #1: US is difficult to acquire

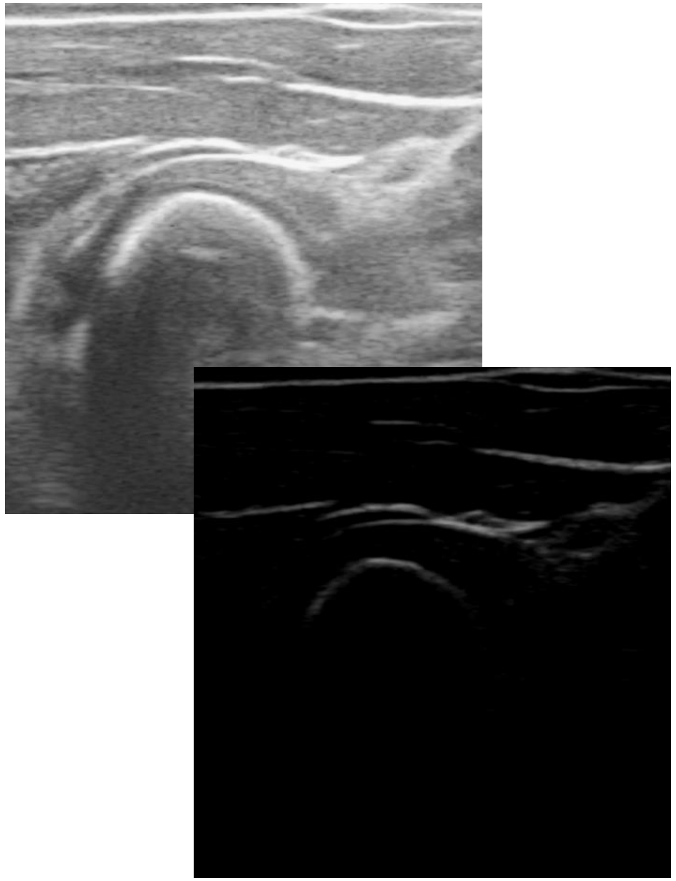

Problem #1: US is difficult to acquire

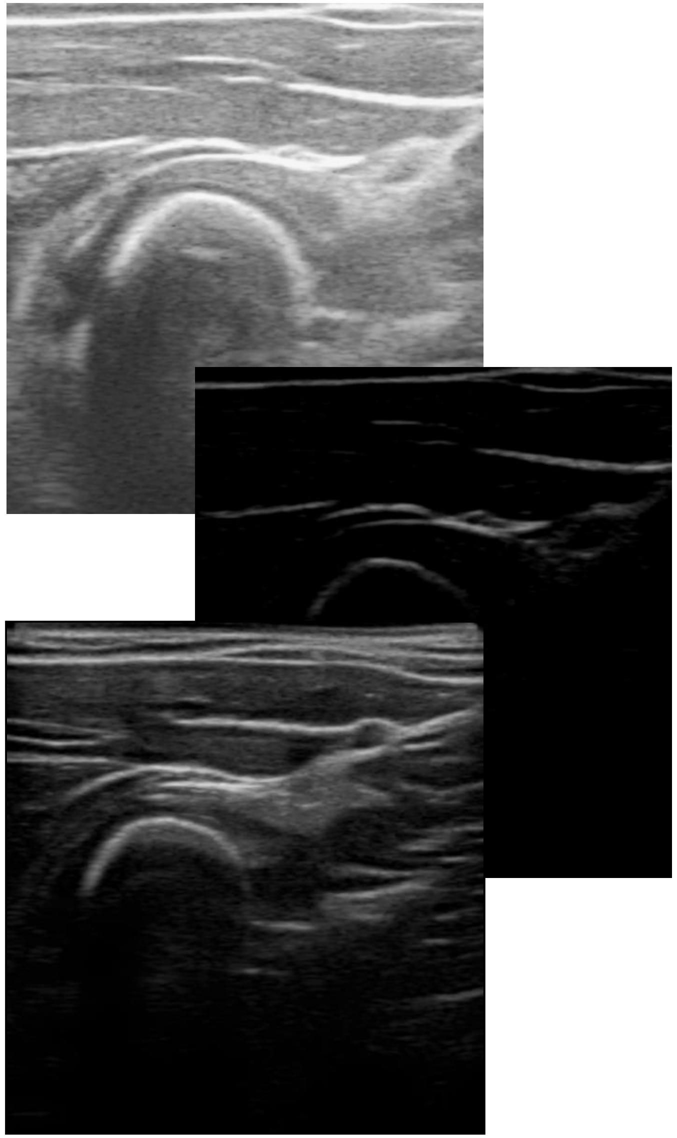

Problem #2: US images are hard to read

Problem #2: US images are hard to read

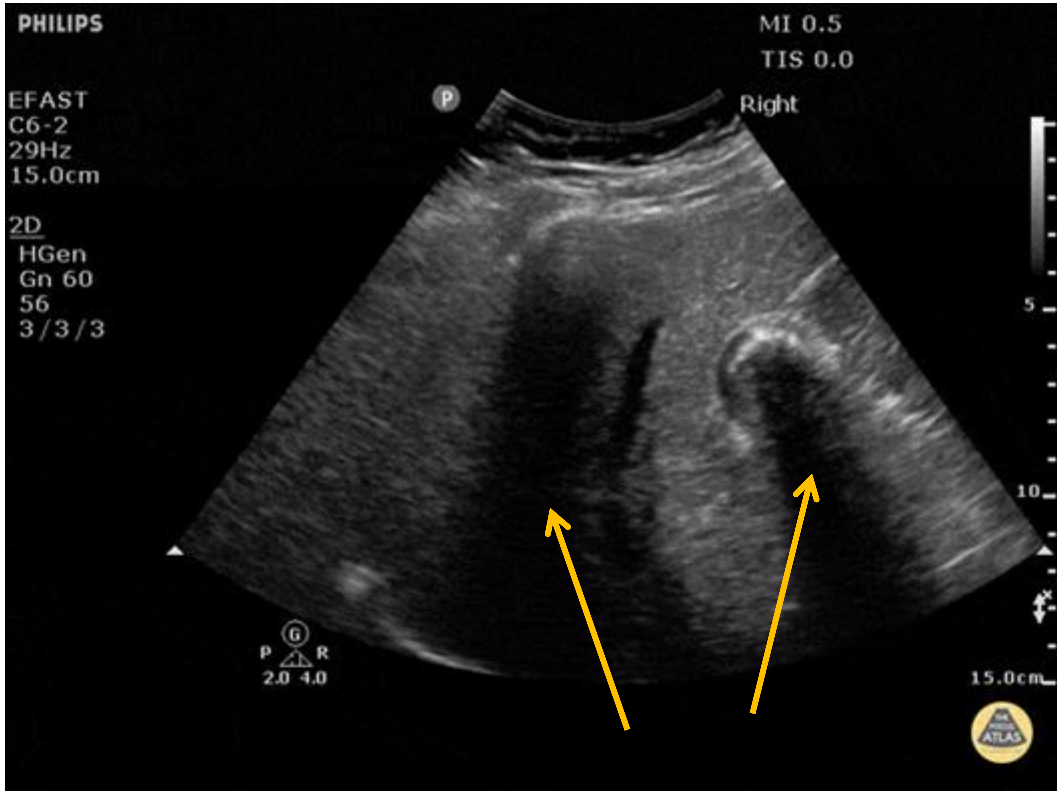

Shadows

Problem #2: US images are hard to read

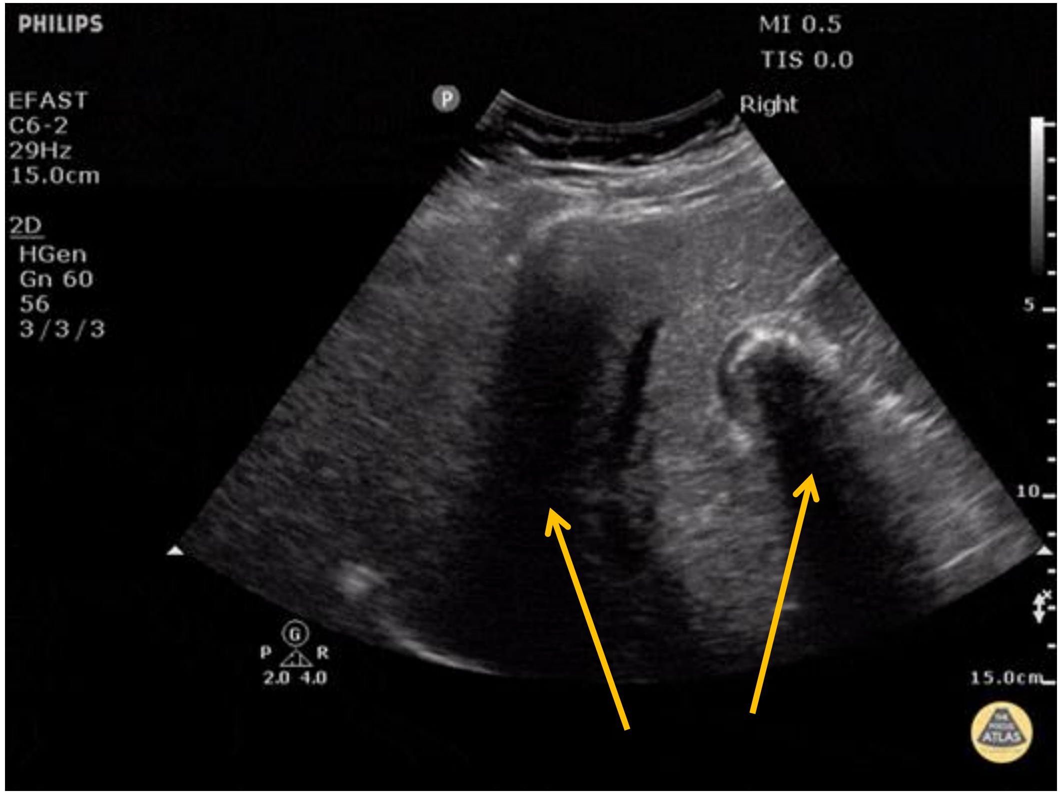

Shadows

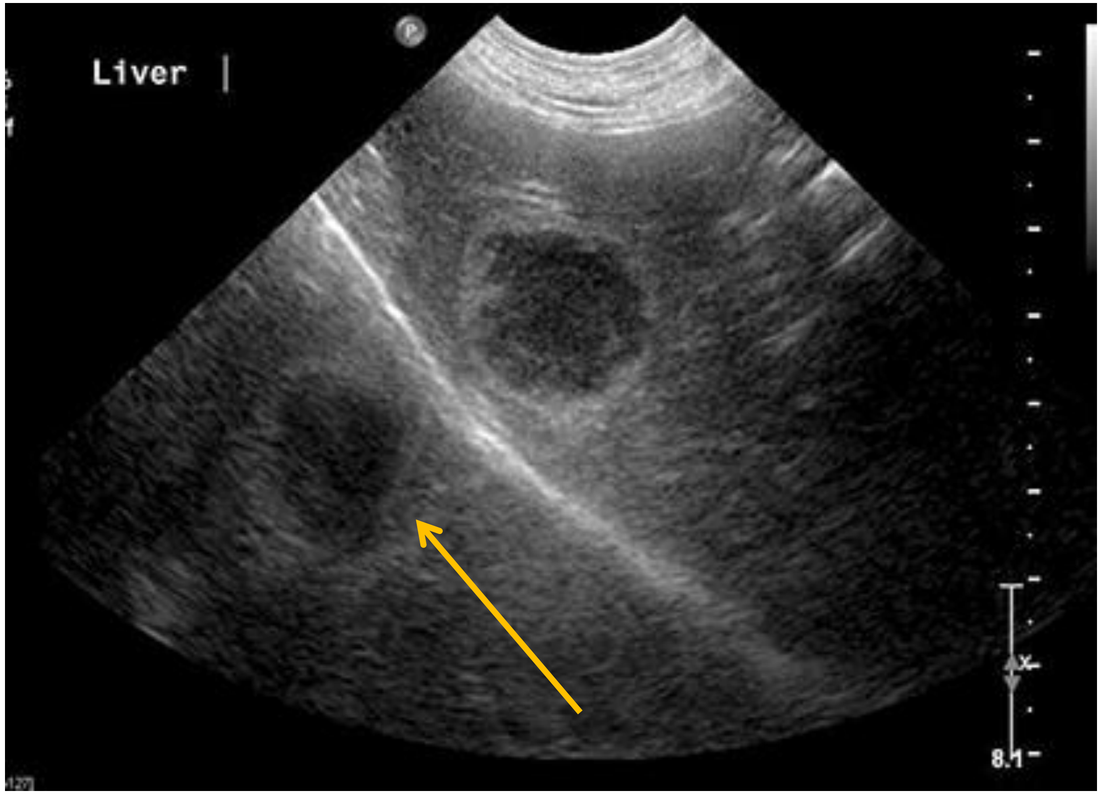

Mirroring

Problem #3: US lack anatomical context

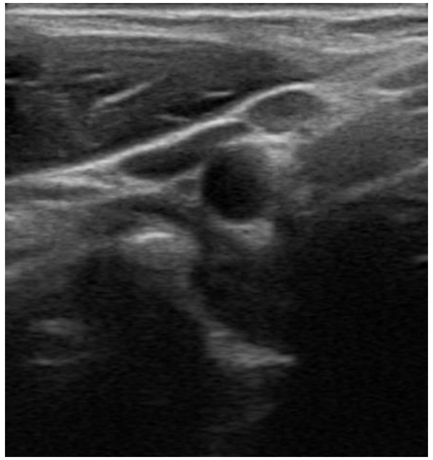

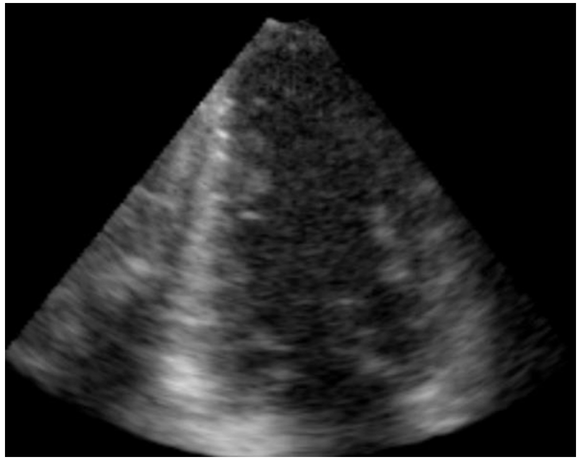

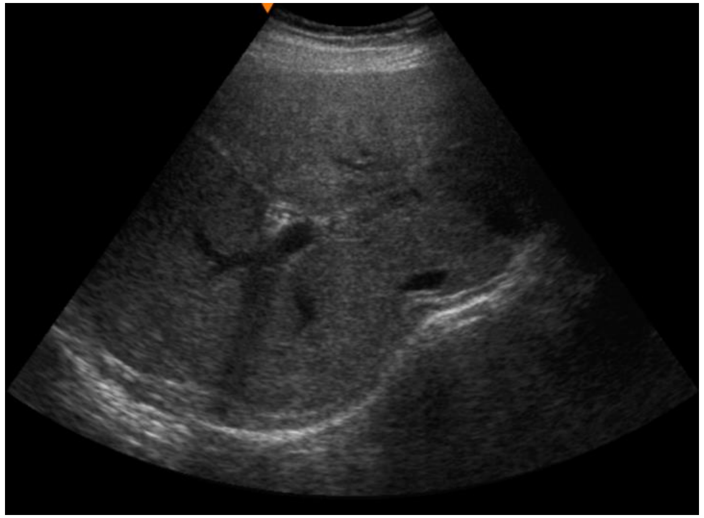

versus

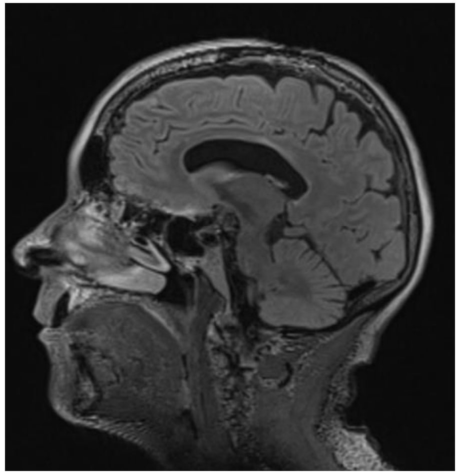

Problem #3: US lack anatomical context

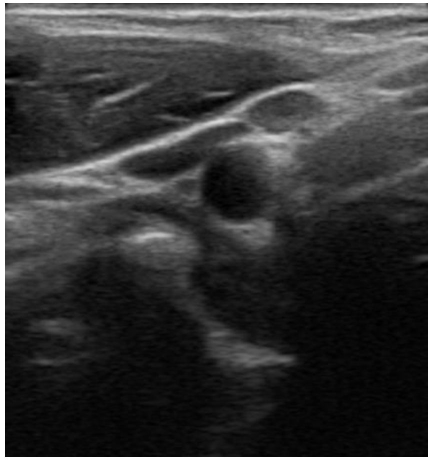

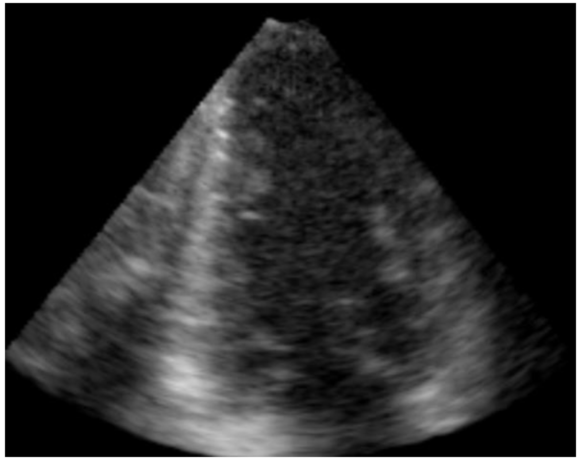

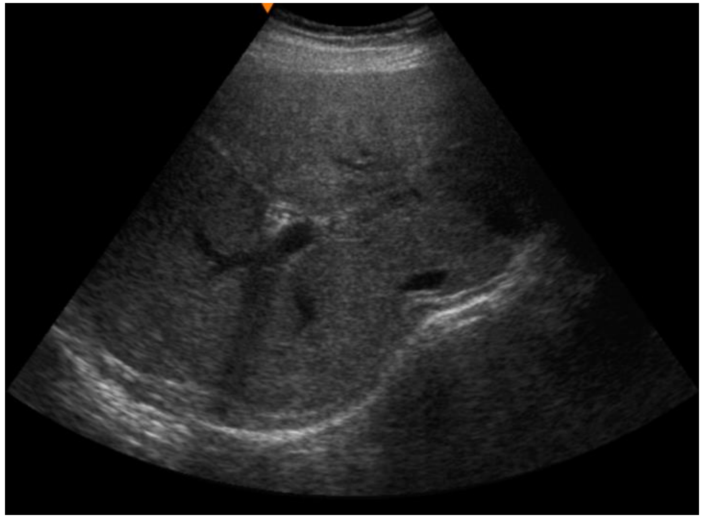

versus

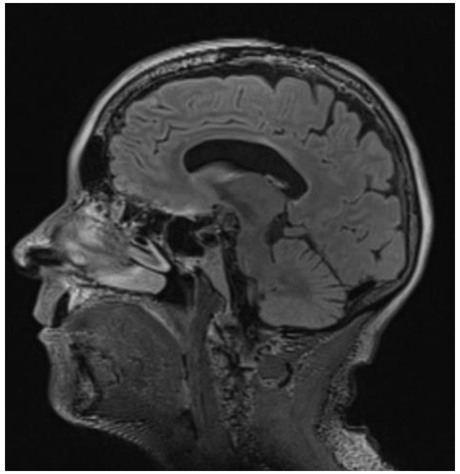

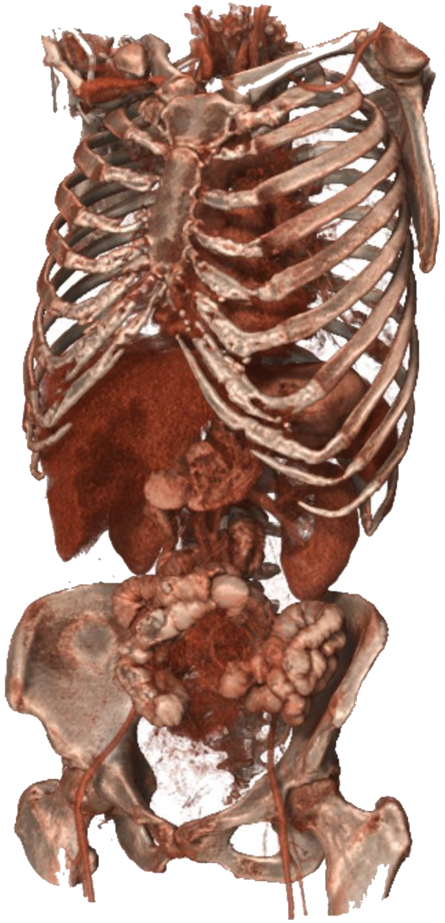

…but ultrasound has a huge potential

✓ Portable

…but ultrasound has a huge potential

✓ Portable ✓ Cheap

5 - 50K $

1M $

…but ultrasound has a huge potential

✓ Portable

✓ Cheap

✓ Safe

5 - 50K $

1M $

…but ultrasound has a huge potential

✓ Portable

5 - 50K $

✓ Cheap

1M $

✓ Safe

✓ High spatial resolution

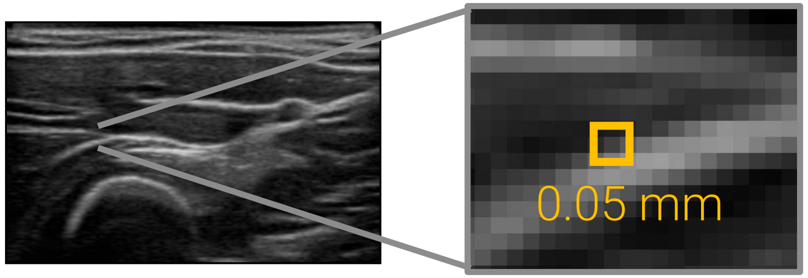

…but ultrasound has a huge potential

✓ Portable

5 - 50K $

✓ Cheap

1M $

✓ Safe

✓ Real-time acquisition
→ suitable for OR

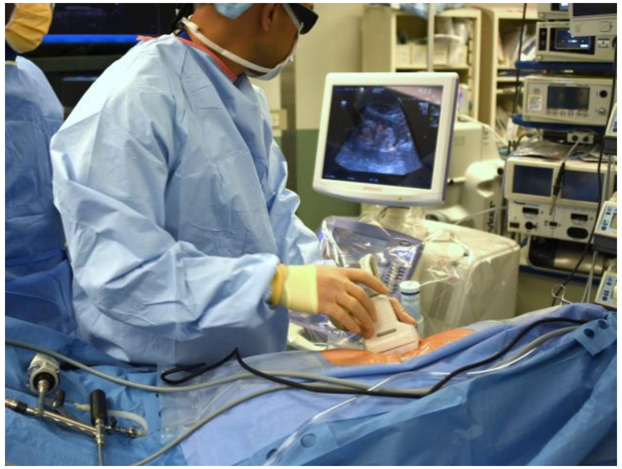

✓ High spatial resolution

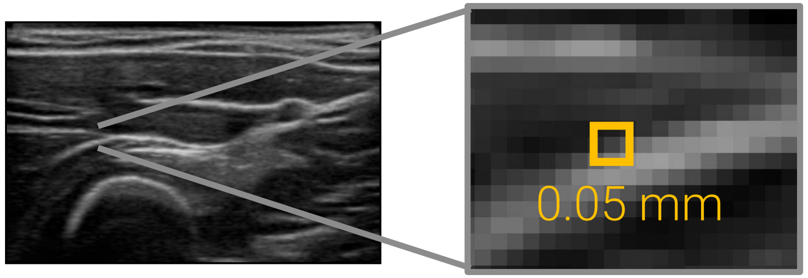

...but ultrasound has a huge potential

✓ Portable

5 - 50K $

✓ Cheap

1M $

✓ Safe

✓ Real-time acquisition
→ suitable for OR

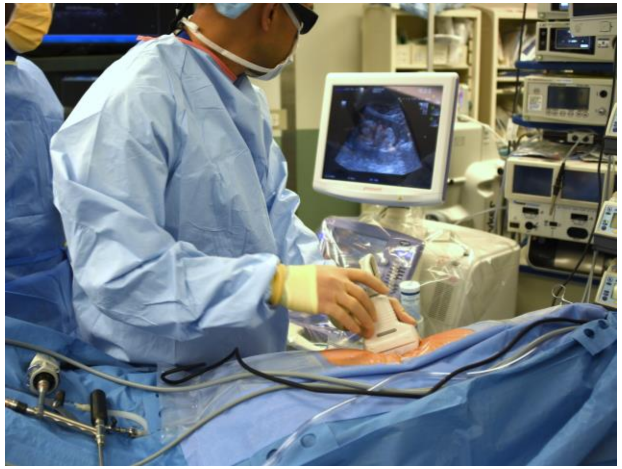

✓ High spatial resolution

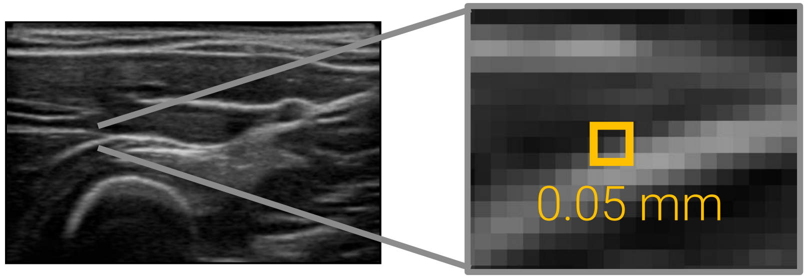

Our vision AI and GPU computing to unlock this potential

Real-Time 2D Image Analysis

Image Filtering

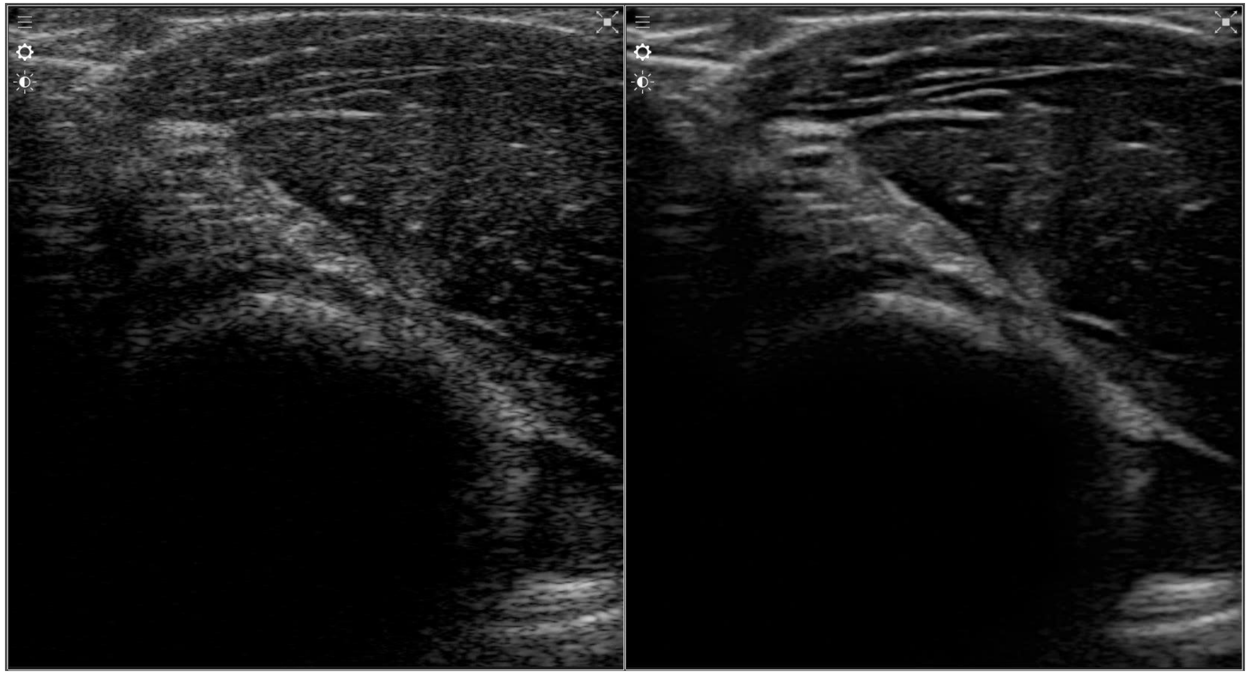

Image Segmentation

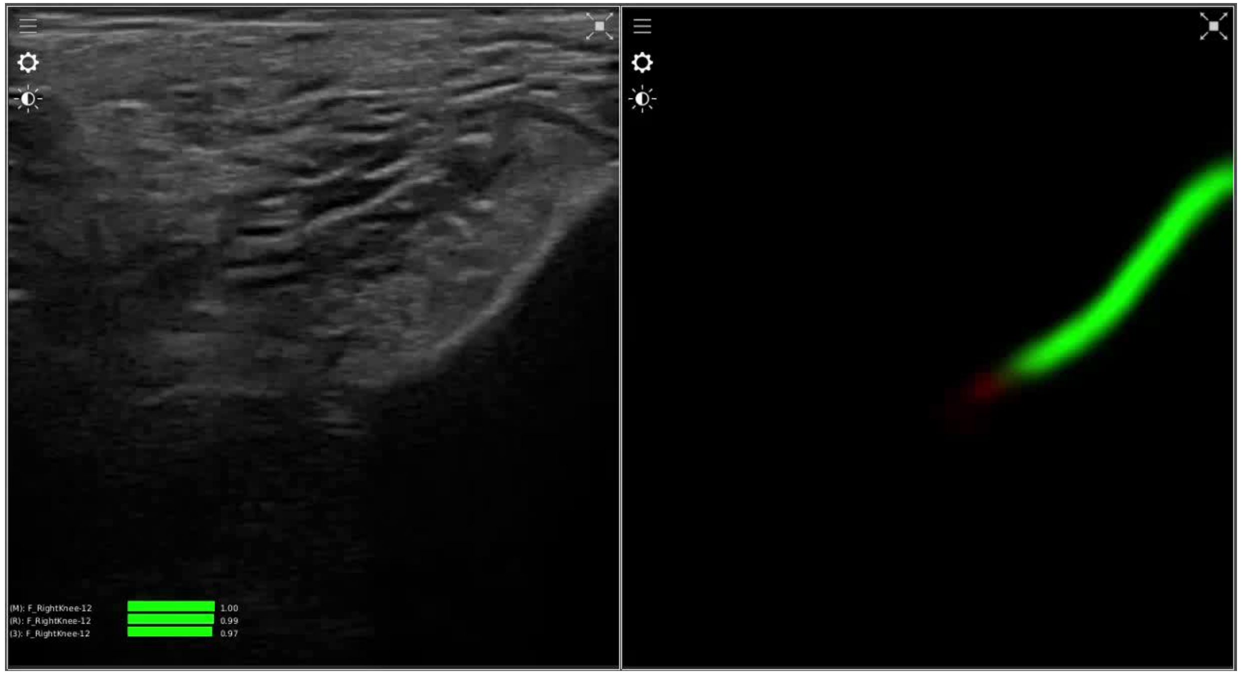

Real-Time 2D Image Analysis

Image Filtering

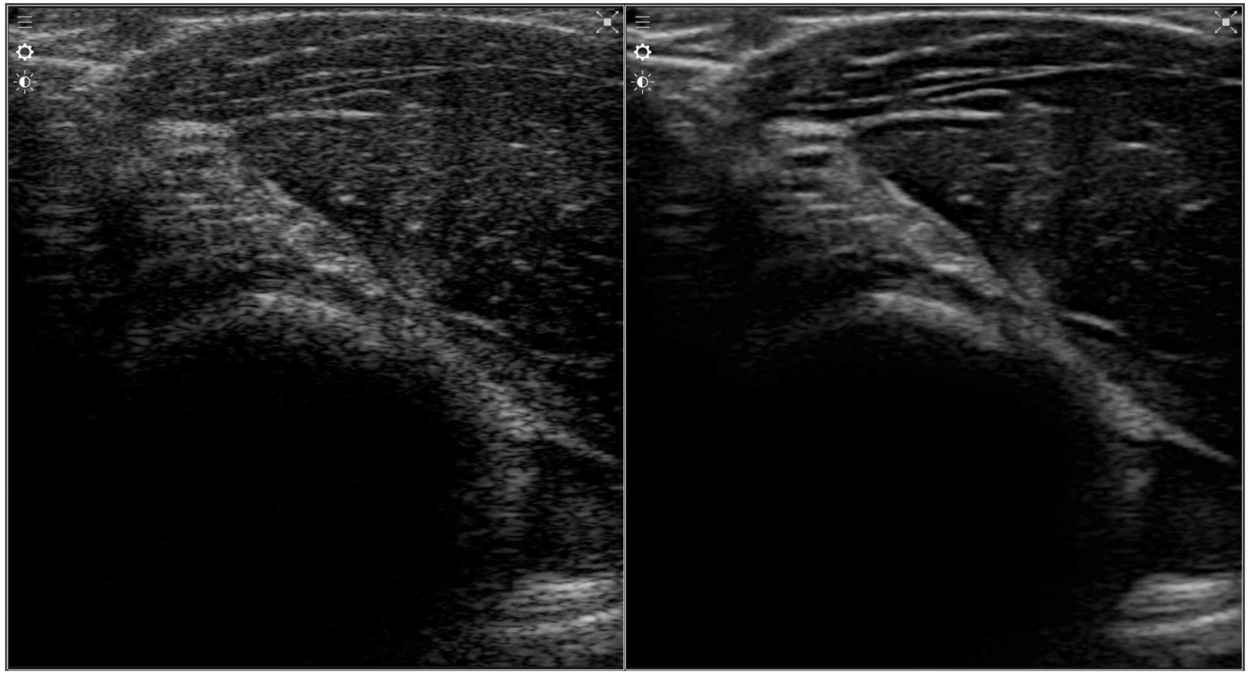

Image Segmentation

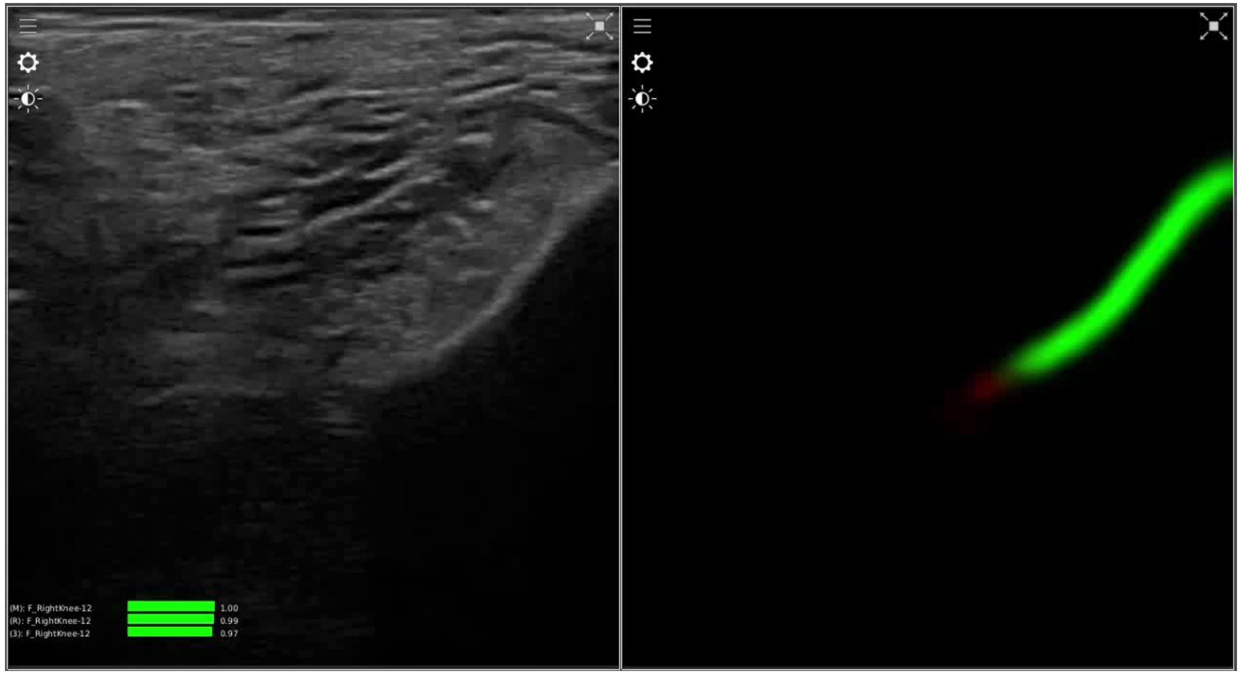

For many clinical applications, we need 3D information (measurements, navigation during surgery, etc.)

From 2D to 3D: Hardware Solutions

Motorized Transducer “wobbler”

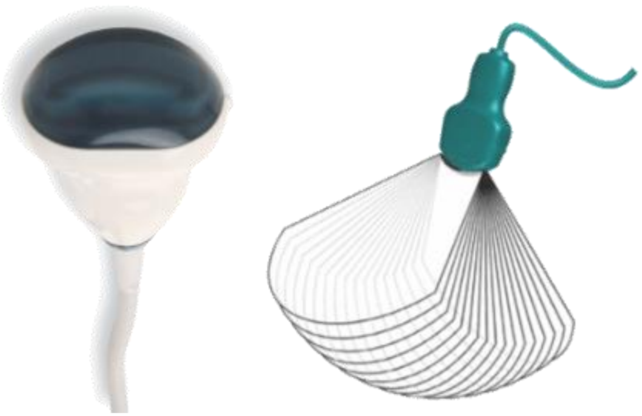

Matrix Array “3D probe”

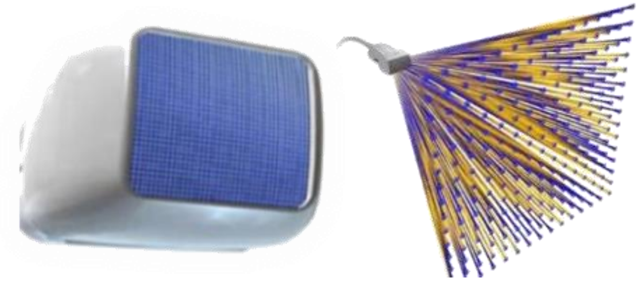

Philips xMatrix

Tracking (optical or electro-magnetic)

Tracked 3D Ultrasound Sweeps

Ultrasound Sweep

2D ultrasound frames, each associated with a 4x4 matrix (position + orientation)

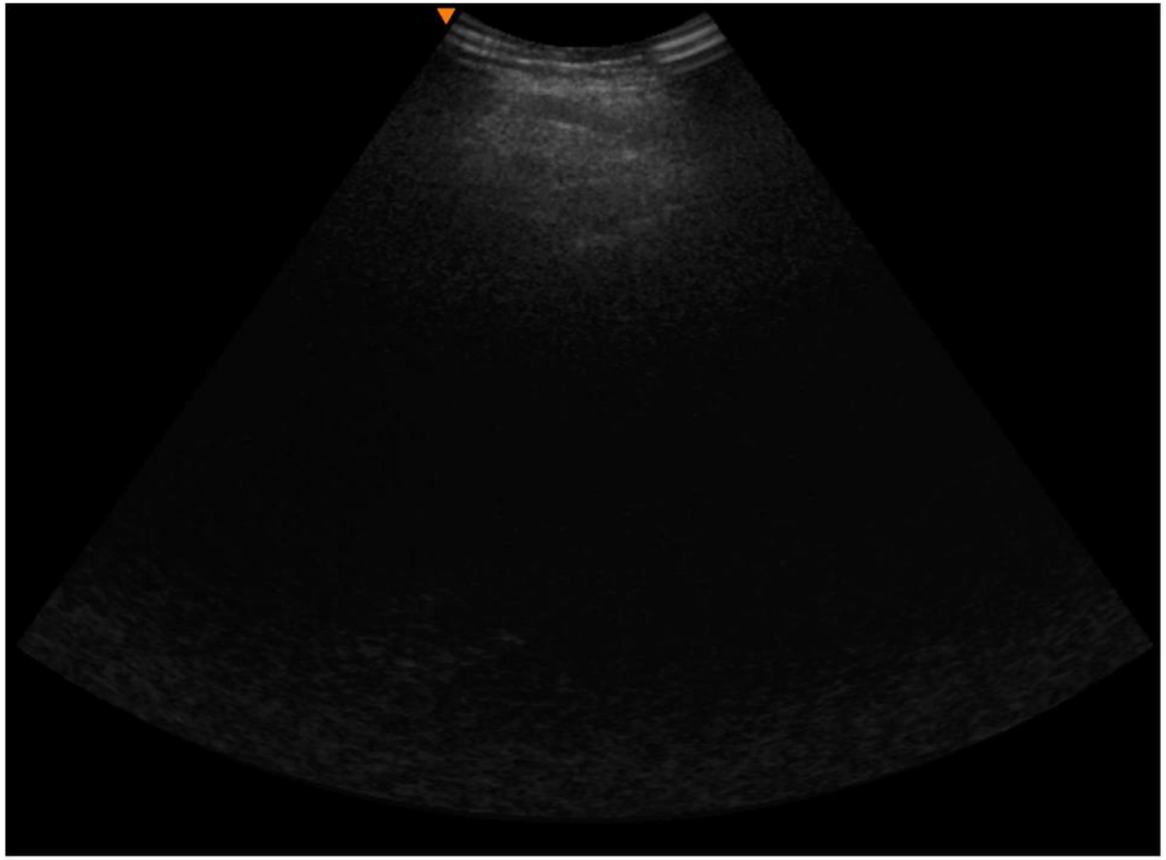

Part 2
Orthopedic Surgery

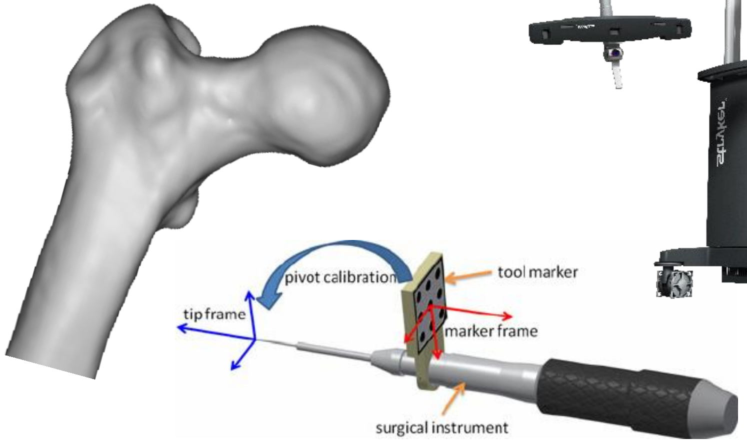

in partnership with

stryker

From planning to navigated surgery

Before Surgery

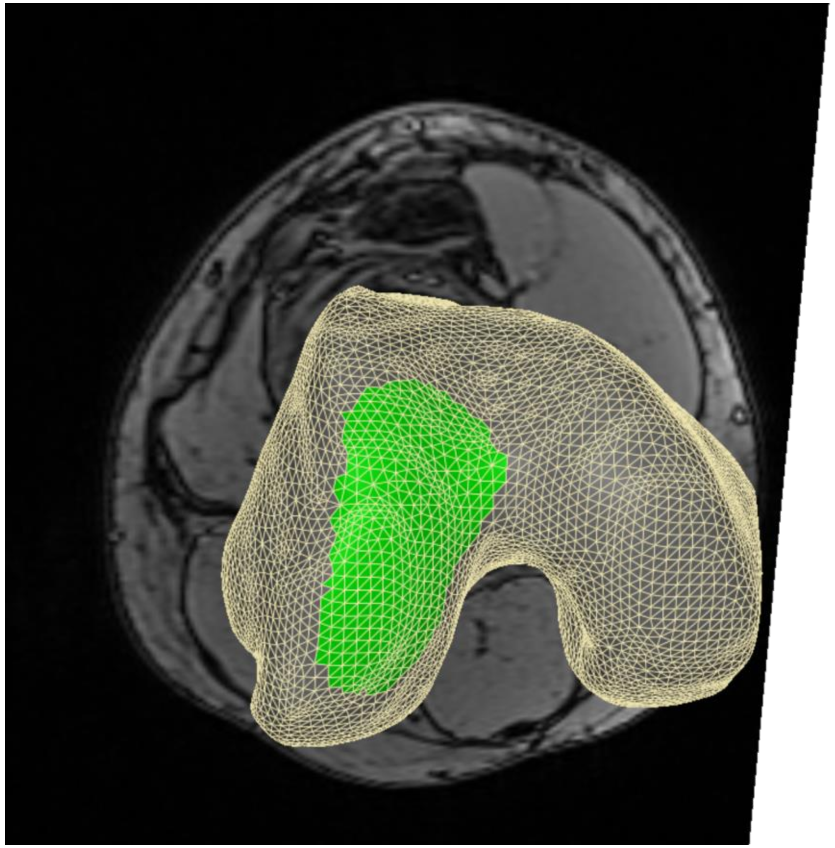

From planning to navigated surgery

Before Surgery

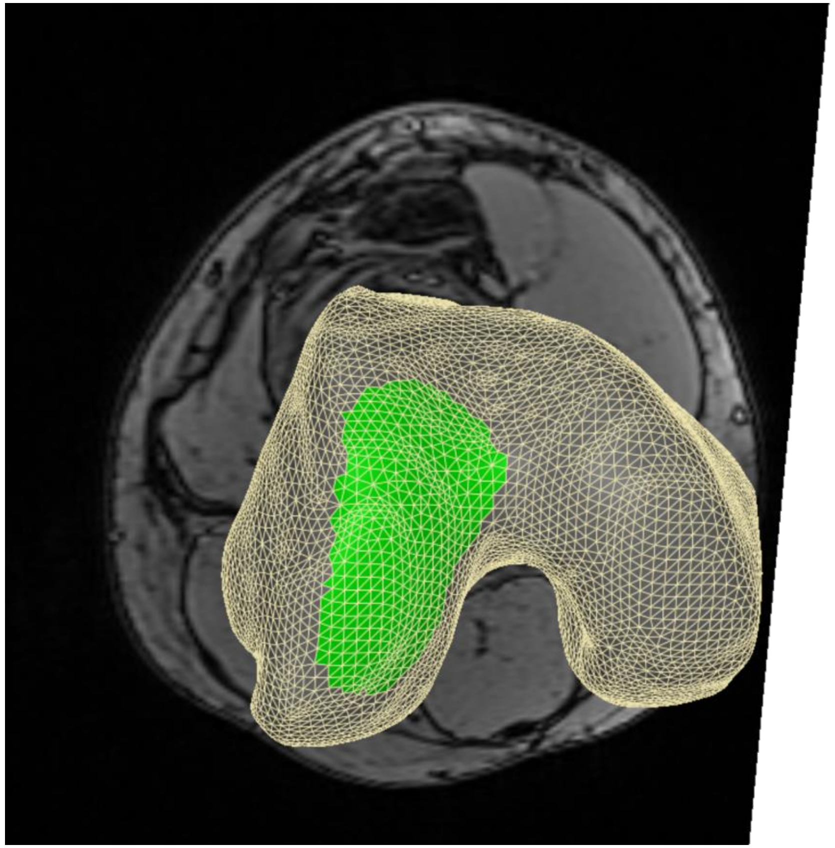

During Surgery

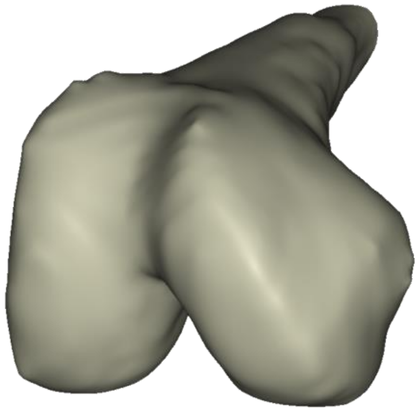

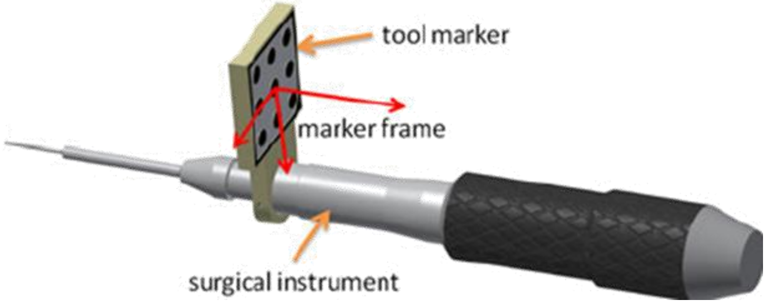

From planning to navigated surgery

Before Surgery

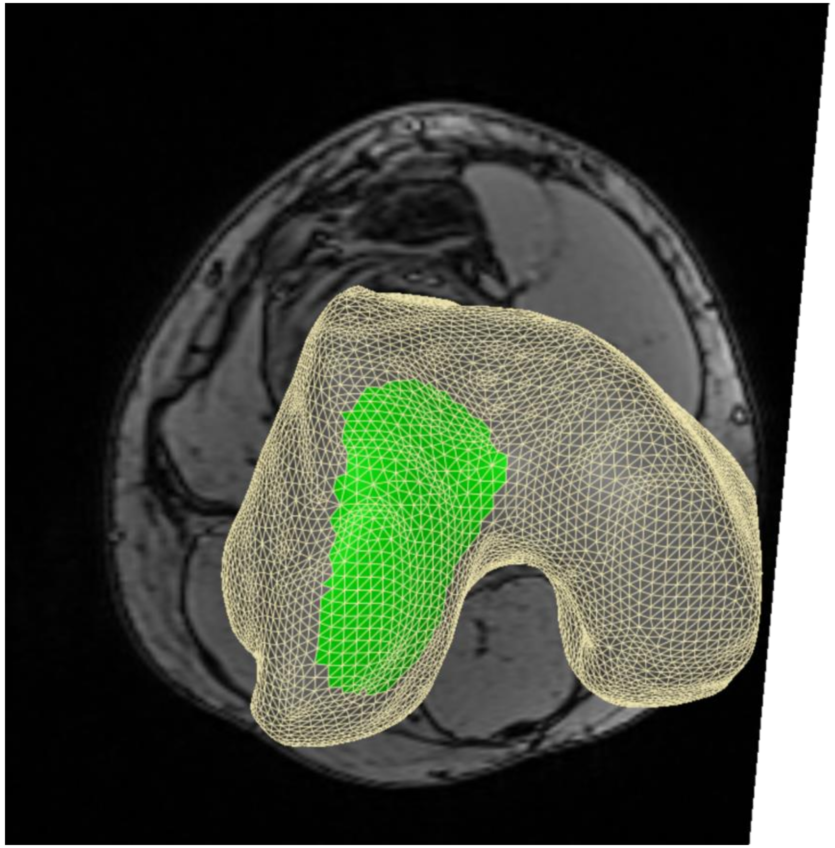

During Surgery

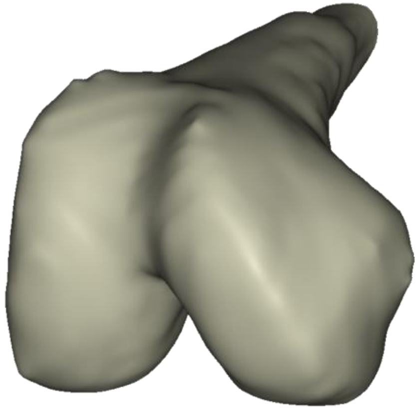

From planning to navigated surgery

Before Surgery

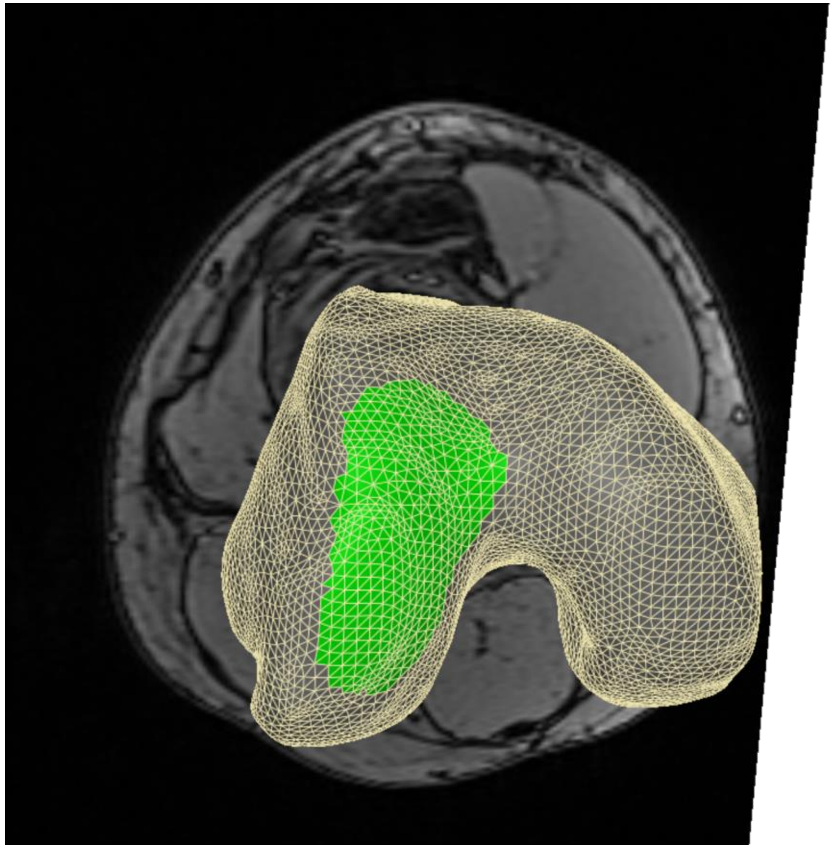

During Surgery

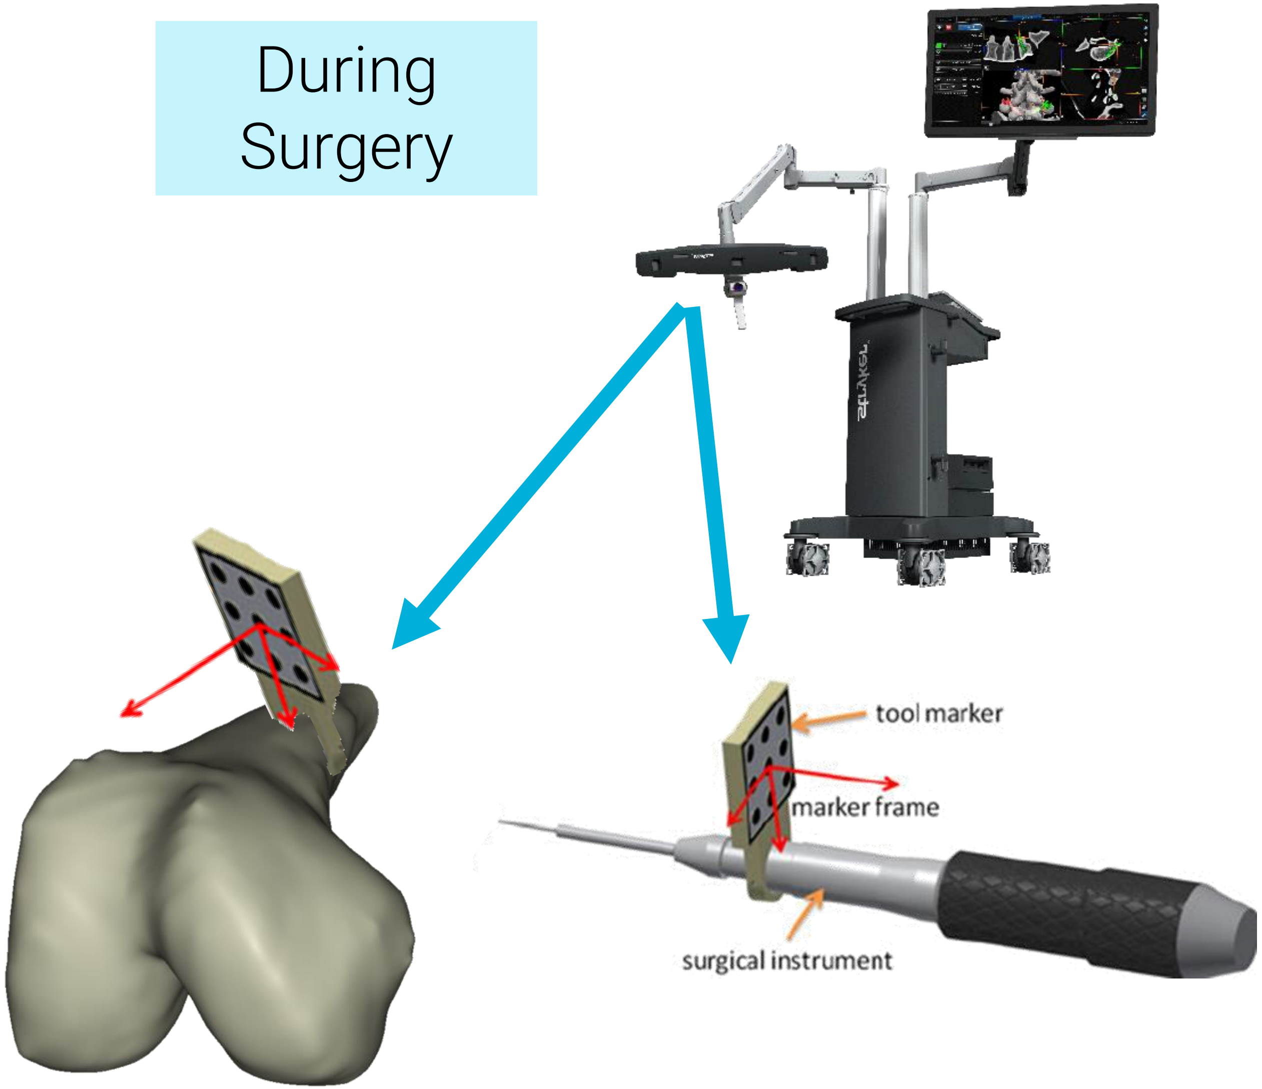

From planning to navigated surgery

Before Surgery

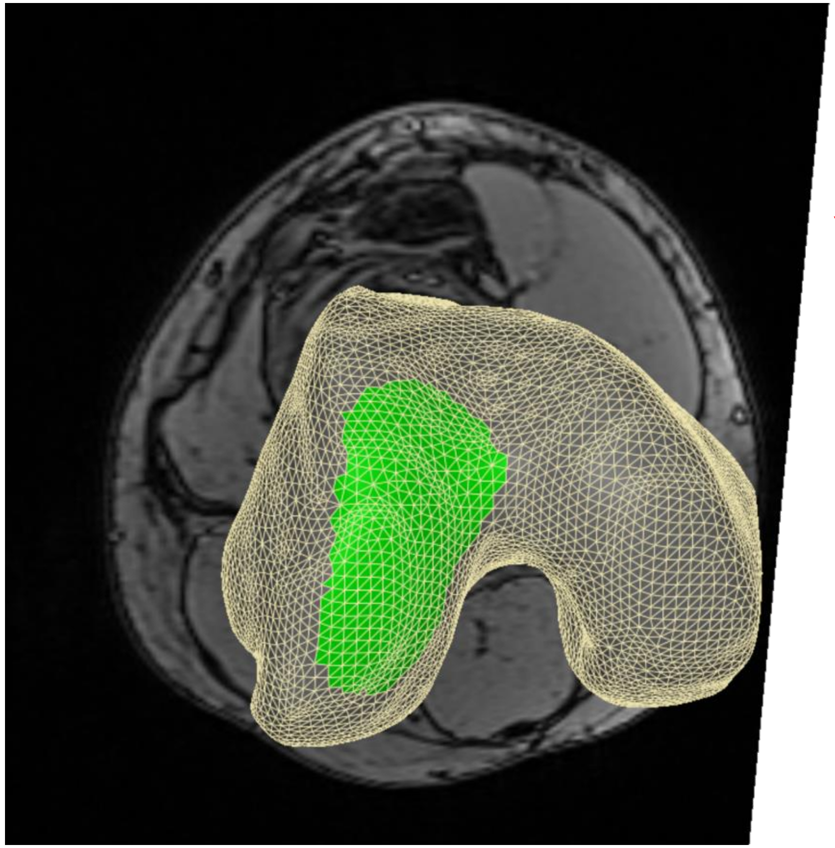

During Surgery

Missing transformation

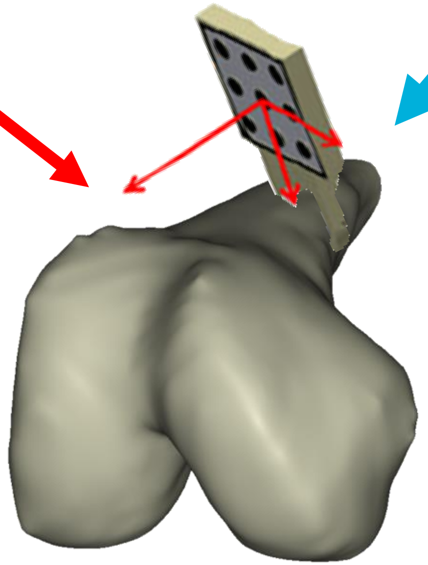

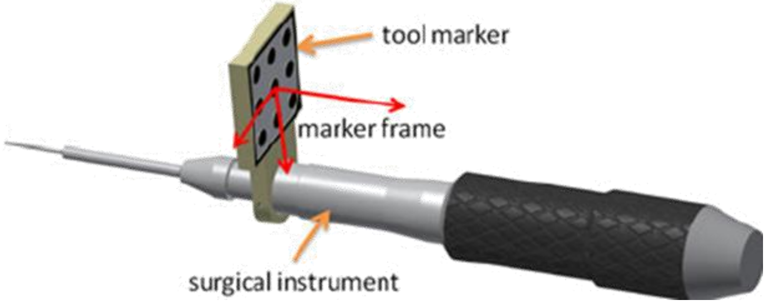

Current workflow

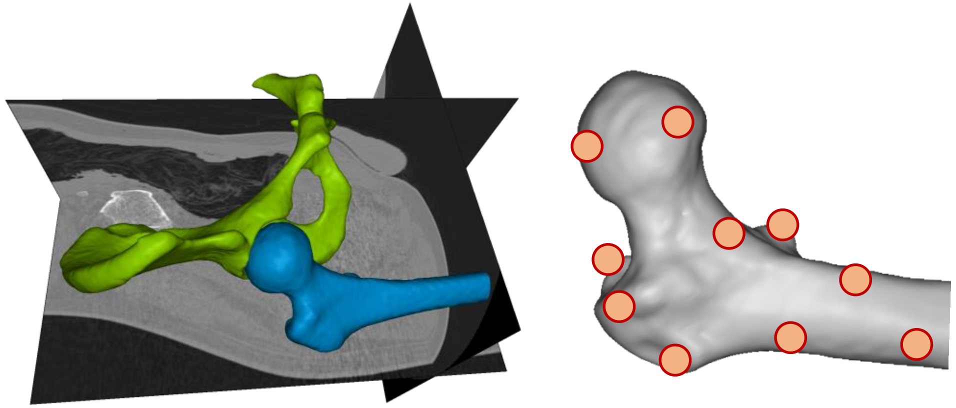

1 Acquire CT/MR image before the operation
Segment the bones and detect landmarks

2 Open the region of surgery
Digitize landmarks on the patient's bone

3 Register pre-op/intra-op landmarks

4 Navigate using the pre-op data

Current workflow

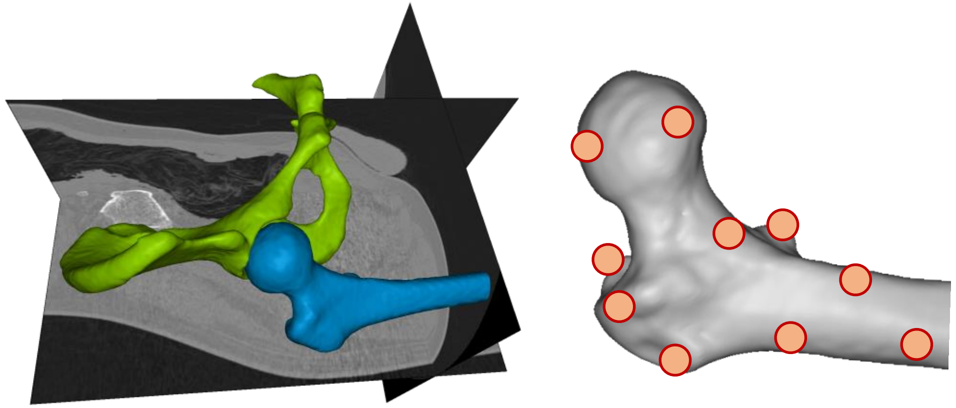

1 Acquire CT/MR image before the operation
Segment the bones and detect landmarks

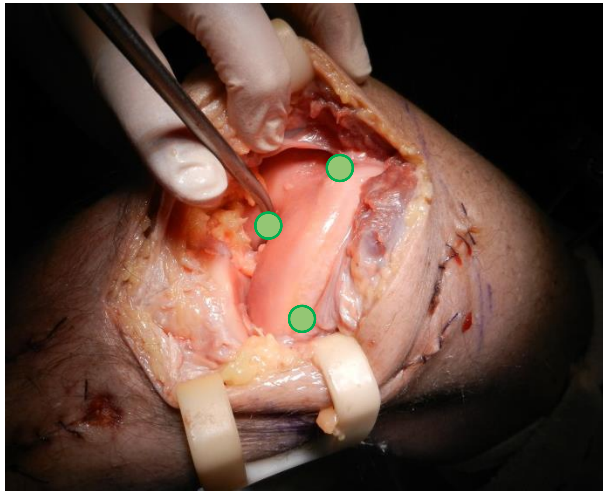

2 Open the region of surgery
Digitize landmarks on the patient's bone

3 Register pre-op/intra-op landmarks

4 Navigate using the pre-op data

Current workflow

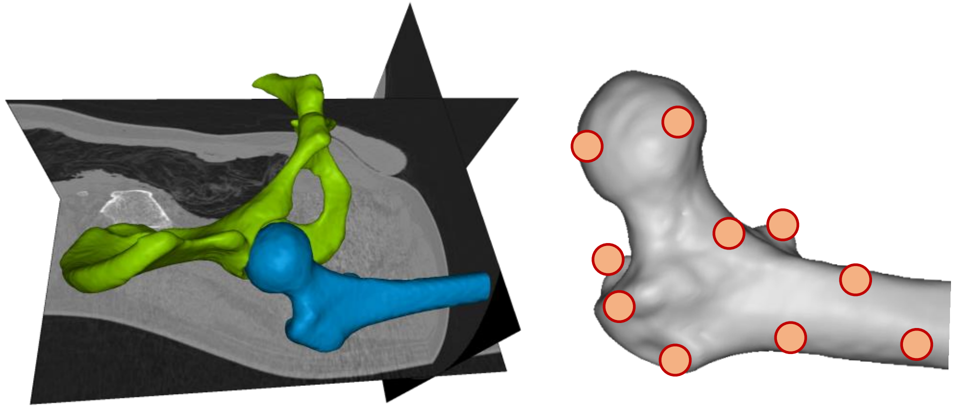

1 Acquire CT/MR image before the operation
Segment the bones and detect landmarks

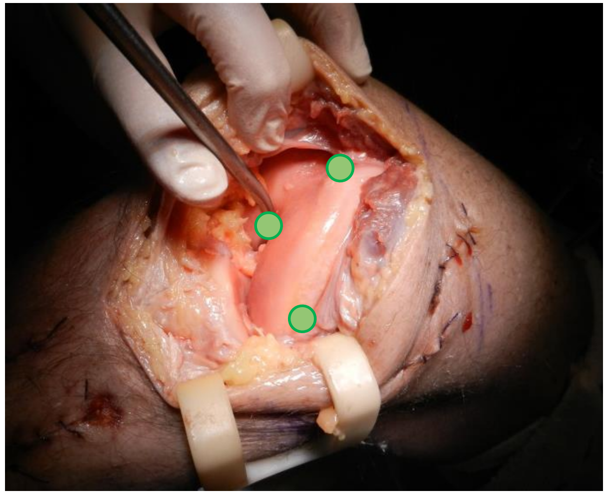

2 Open the region of surgery
Digitize landmarks on the patient's bone

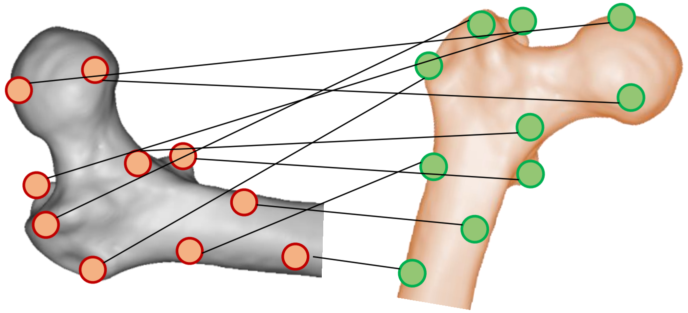

3 Register pre-op/intra-op landmarks

4 Navigate using the pre-op data

Current workflow

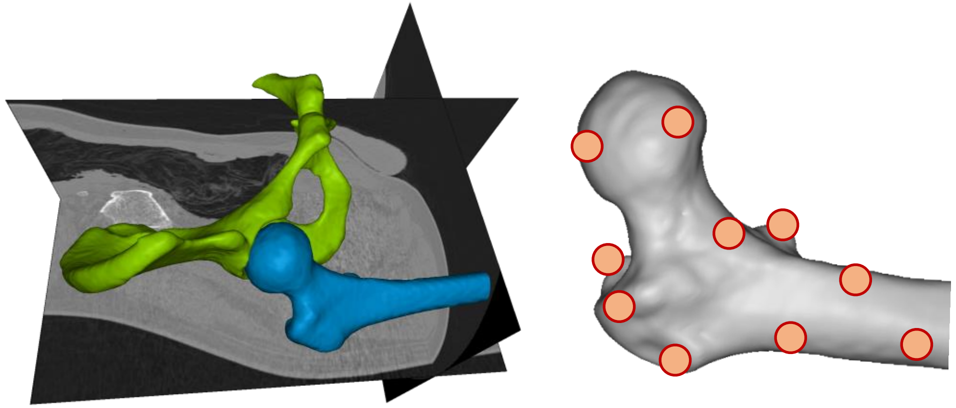

1 Acquire CT/MR image before the operation
Segment the bones and detect landmarks

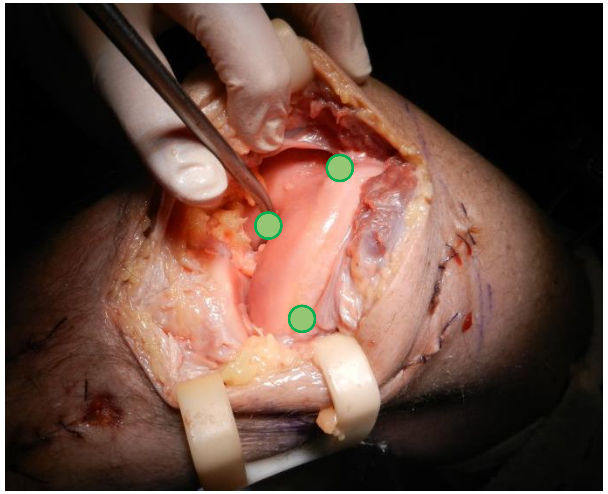

2 Open the region of surgery
Digitize landmarks on the patient's bone

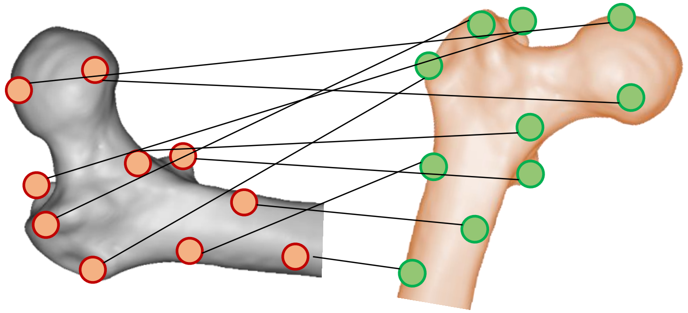

3 Register pre-op/intra-op landmarks

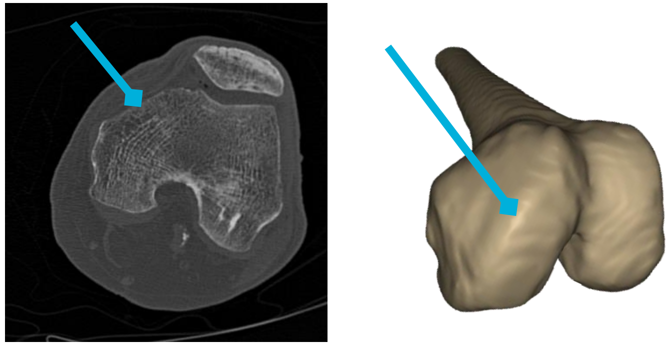

4 Navigate using the pre-op data

Current workflow

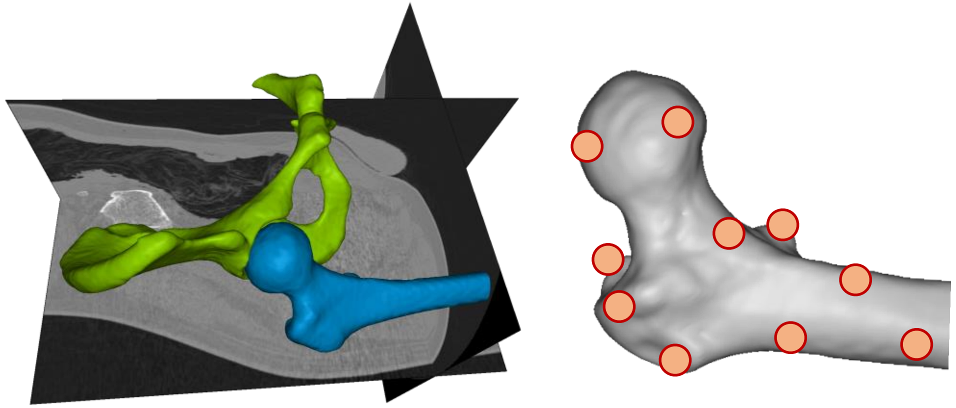

1 Acquire CT/MR image before the operation
Segment the bones and detect landmarks

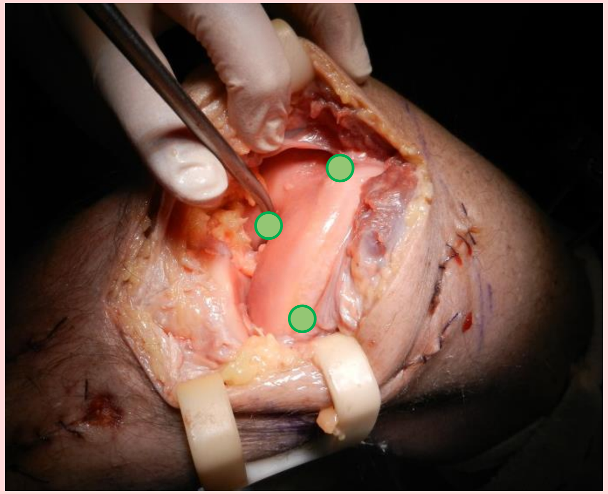

2 Open the region of surgery
Digitize landmarks on the patient's bone

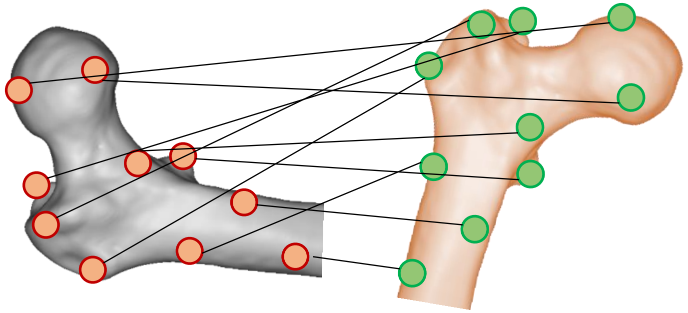

3 Register pre-op/intra-op landmarks

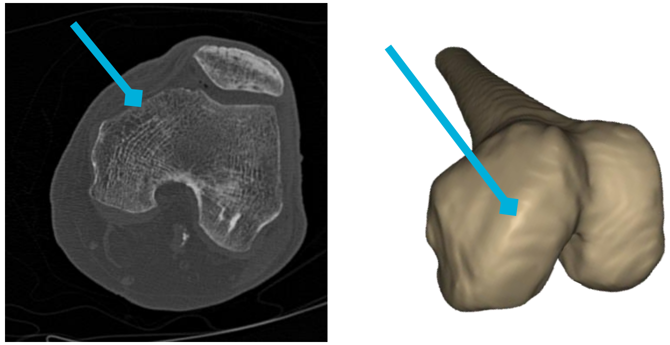

4 Navigate using the pre-op data

Ultrasound-based workflow

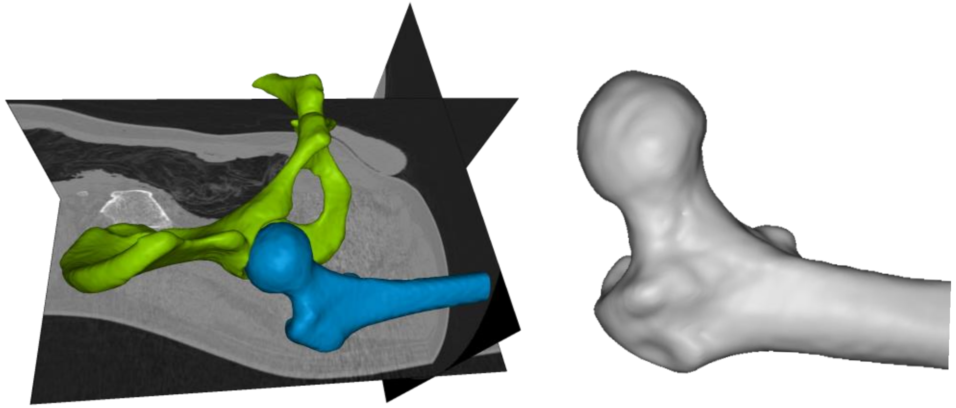

1 Acquire CT/MR image before the operation
Segment the bones

2 Acquire a tracked 3D Ultrasound sweep
Extract the bone surface

3 Register pre-op/intra-op bone surface

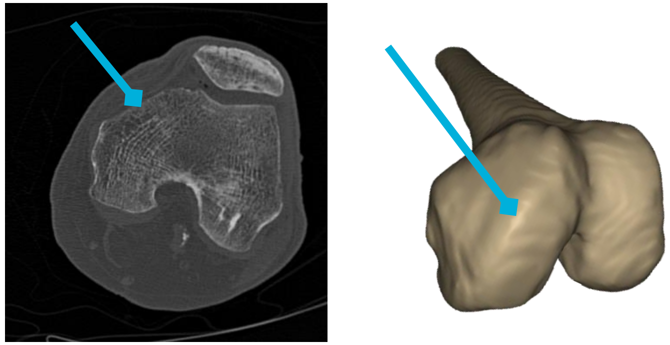

4 Navigate using the pre-op data

Ultrasound-based workflow

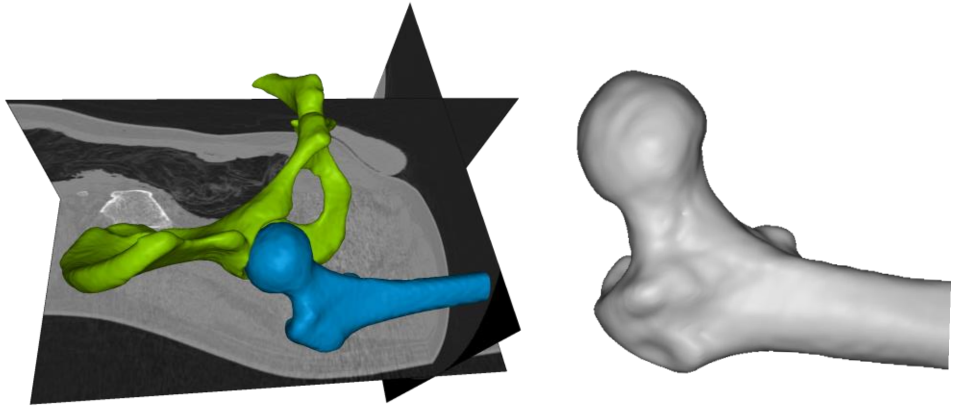

1 Acquire CT/MR image before the operation
Segment the bones

2 Acquire a tracked 3D Ultrasound sweep
Extract the bone surface

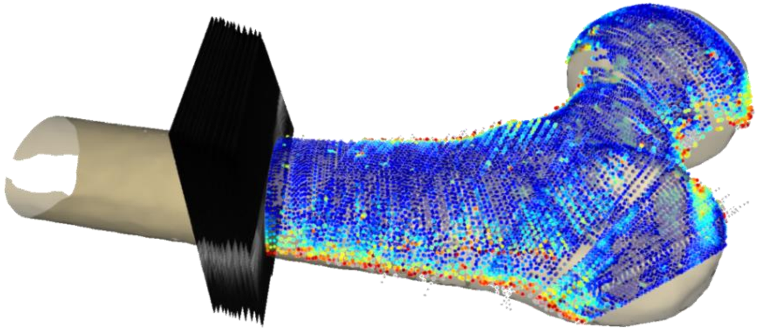

3 Register pre-op/intra-op bone surface

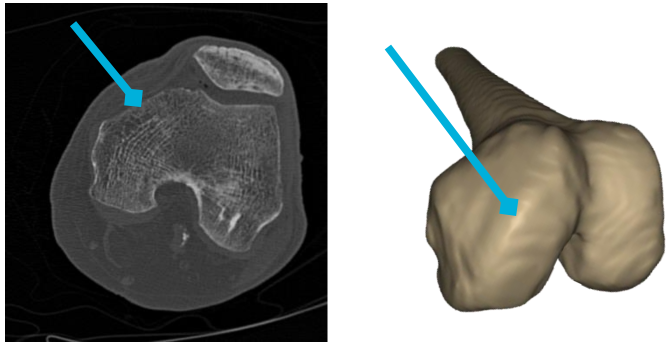

4 Navigate using the pre-op data

Deep Learning for Bone Detection

High shape and appearance variability

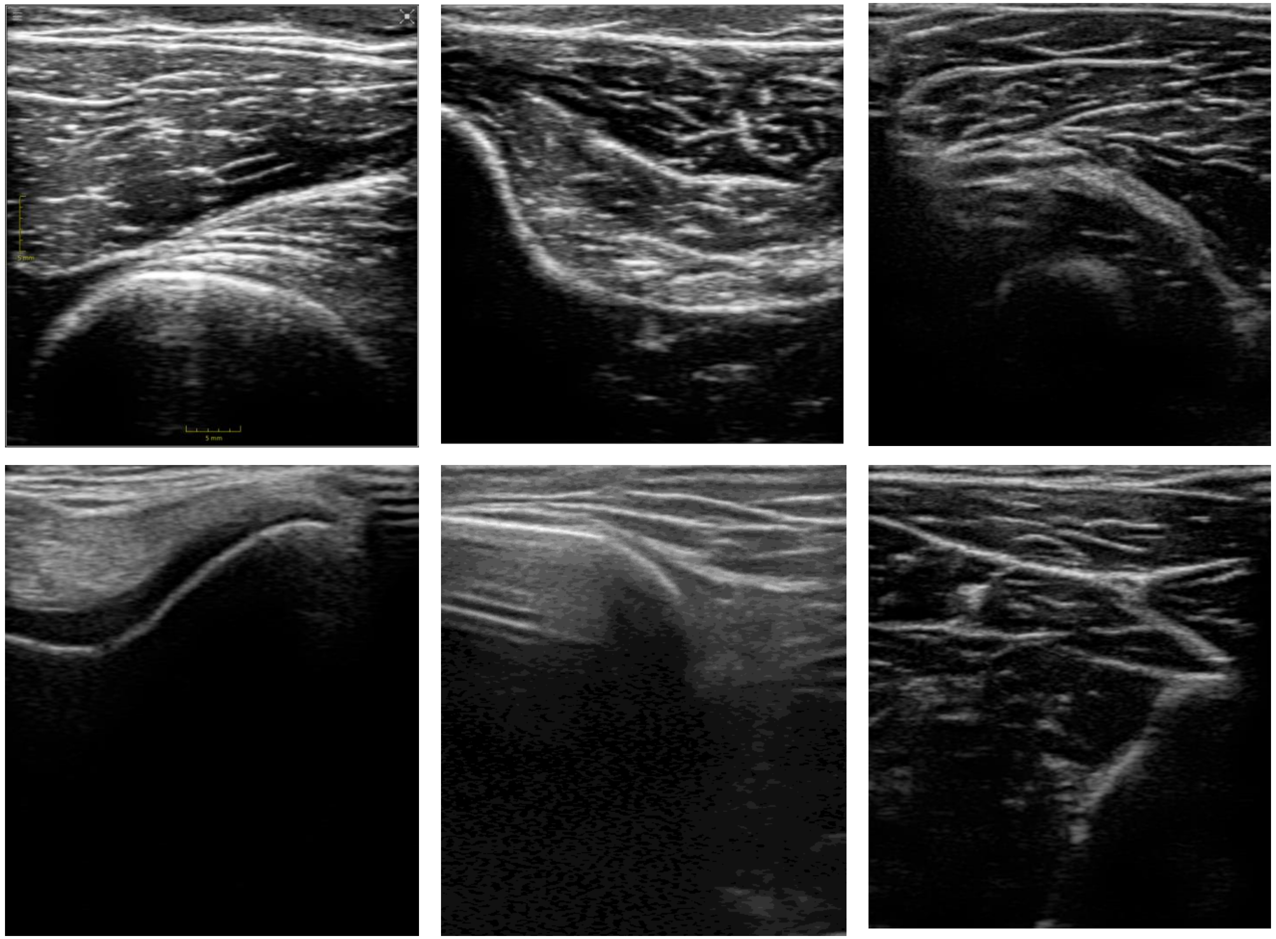

Deep Learning for Bone Detection

High shape and appearance variability

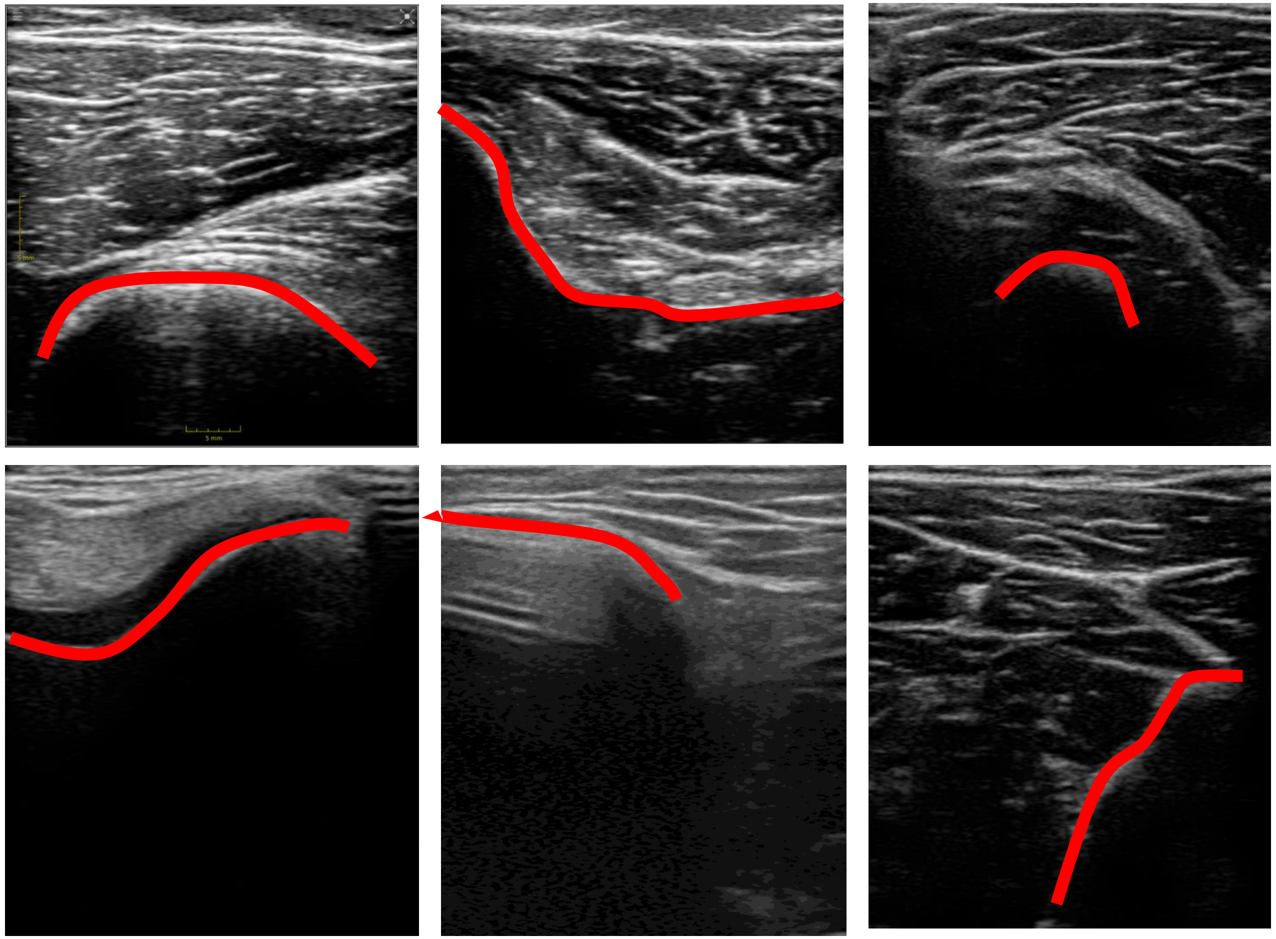

Network Architecture for Segmentation

U-Net Architecture
(most popular for medical image segmentation)

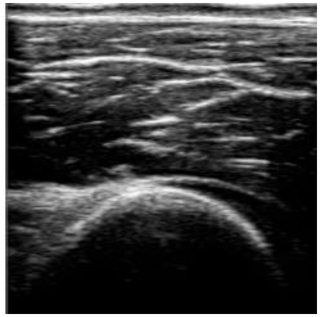

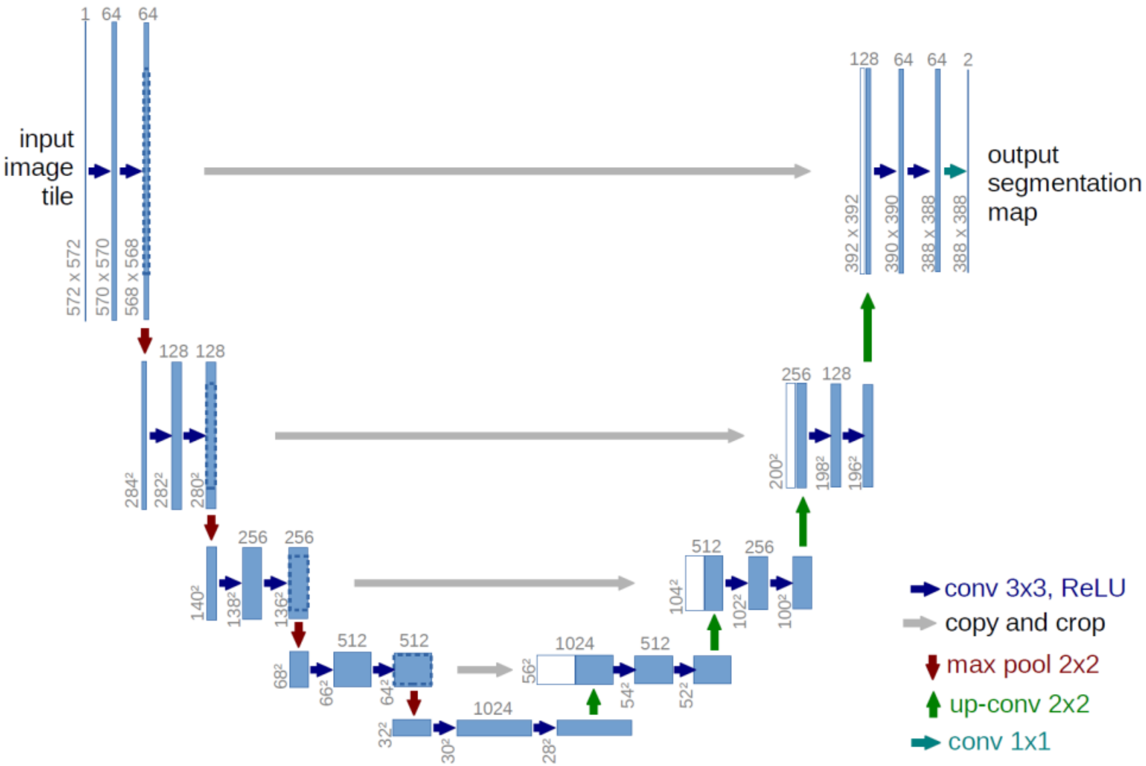

from https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

Bone Segmentation Results

US Image | Random Forest | Neural Network

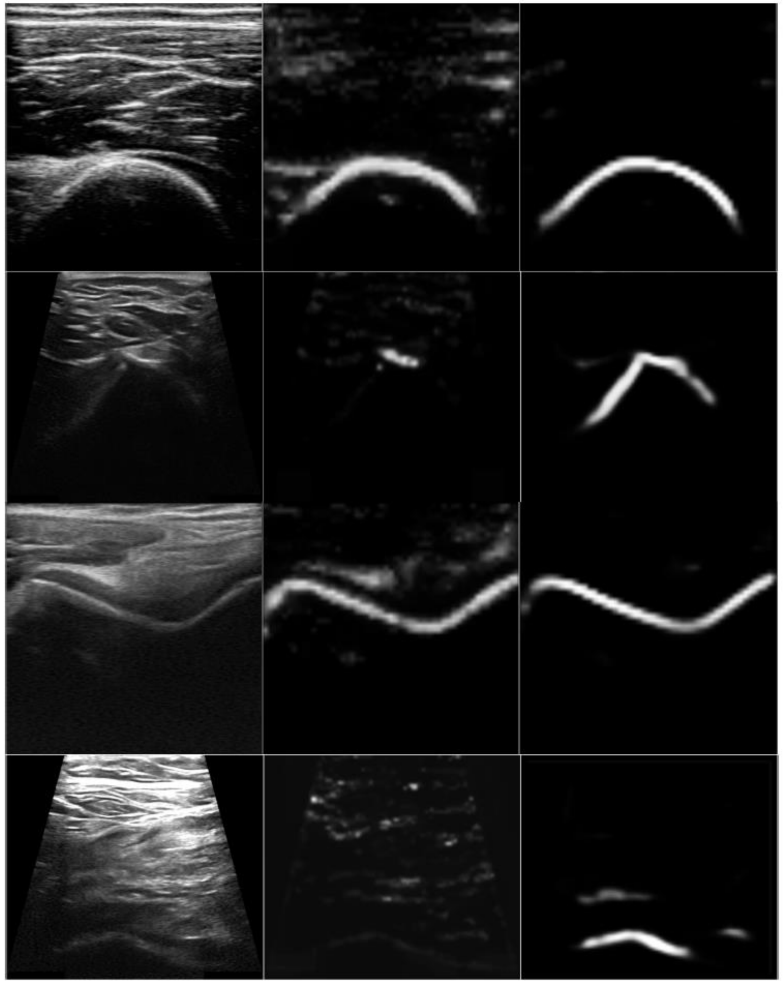

The segmentation is then refined at the pixel-level

Bone Segmentation Results

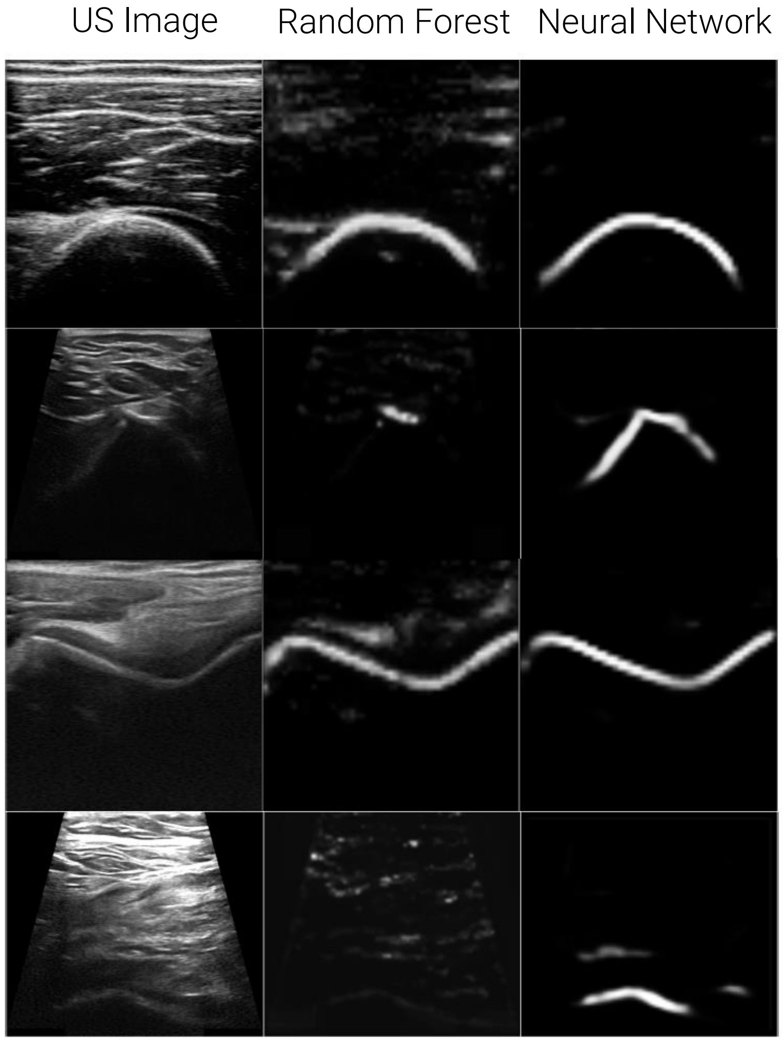

The segmentation is then refined at the pixel-level

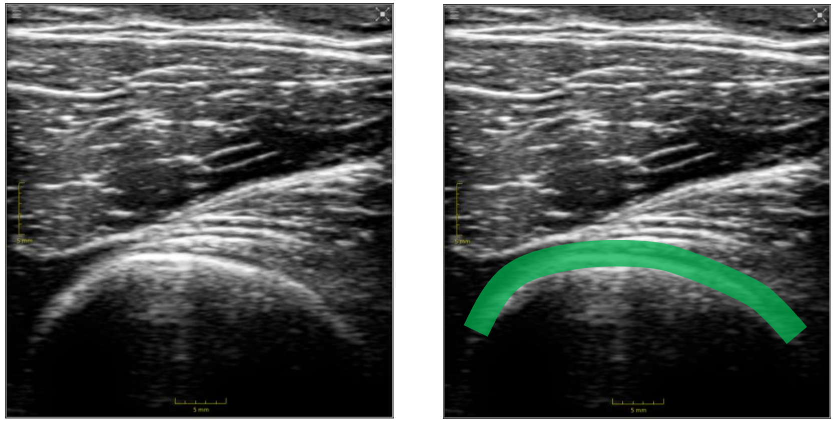

Bone Segmentation Results

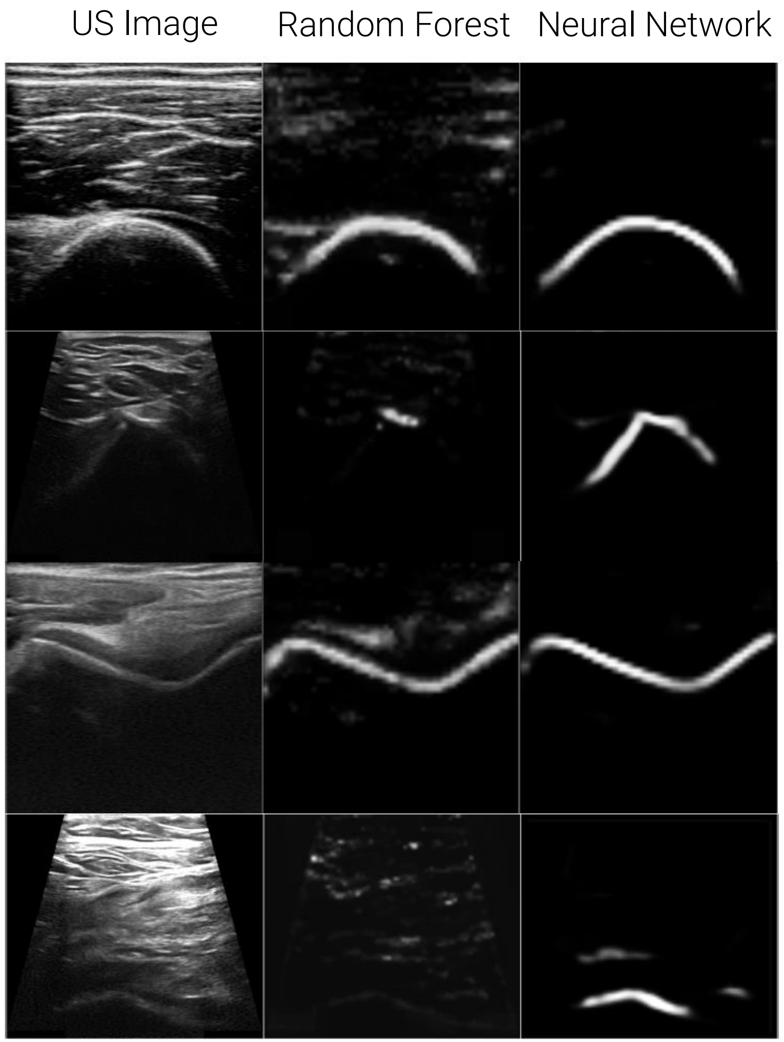

The segmentation is then refined at the pixel-level

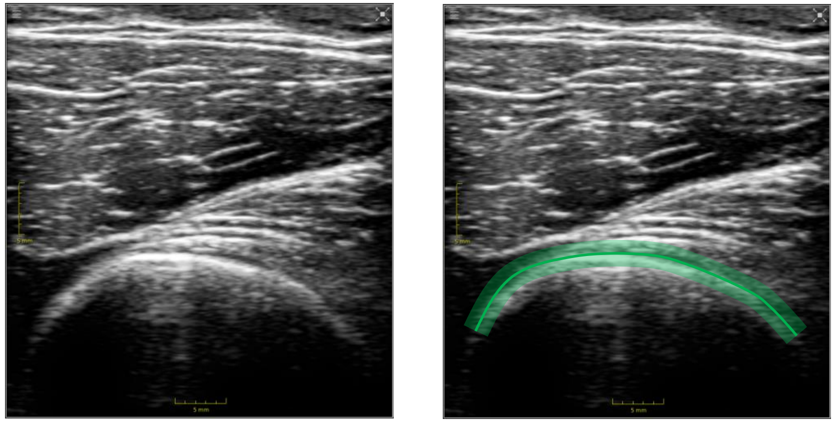

Bone Segmentation Results

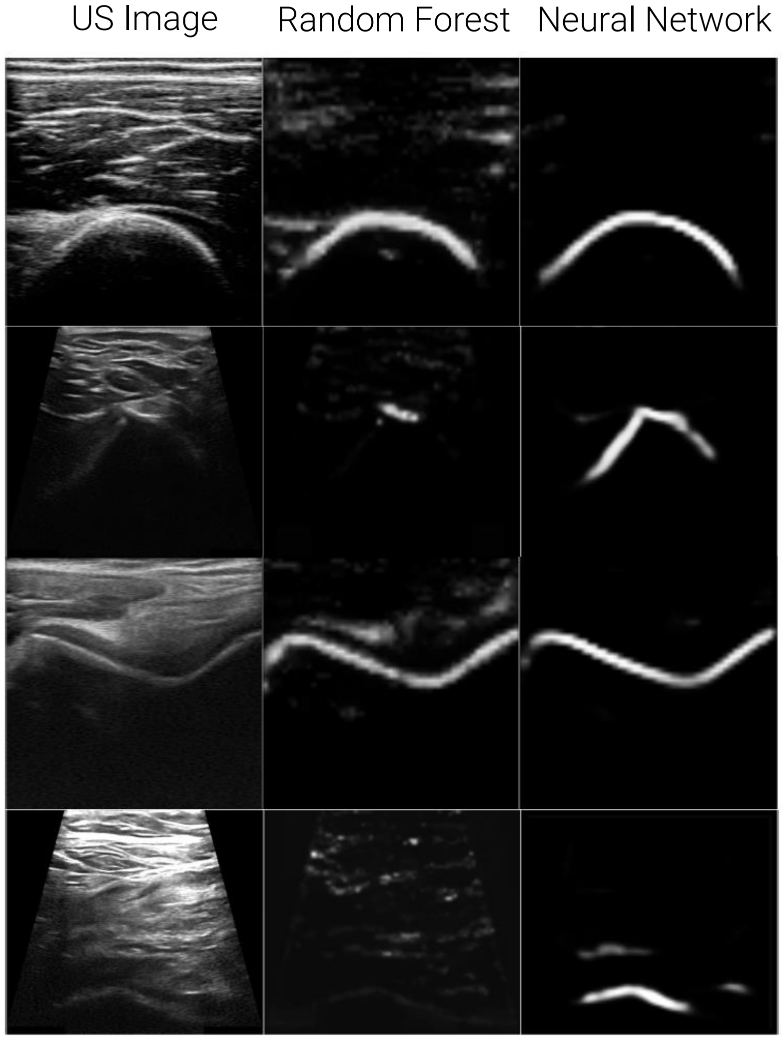

The segmentation is then refined at the pixel-level

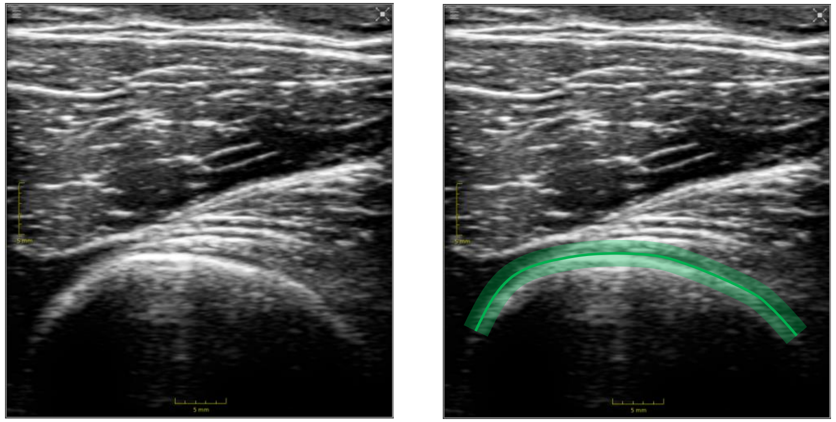

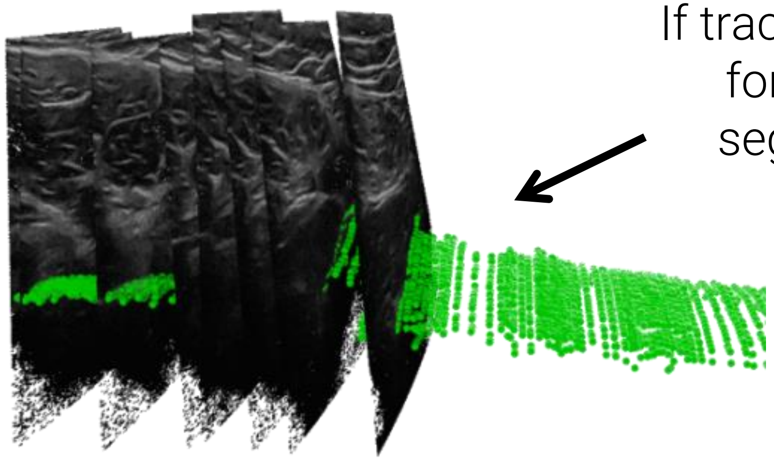

If tracking data is available for each frame, a 3D segmentation can be generated

Point Cloud to Surface Registration

Optimization problem

Minimize the distance between each point and the closest point on the surface

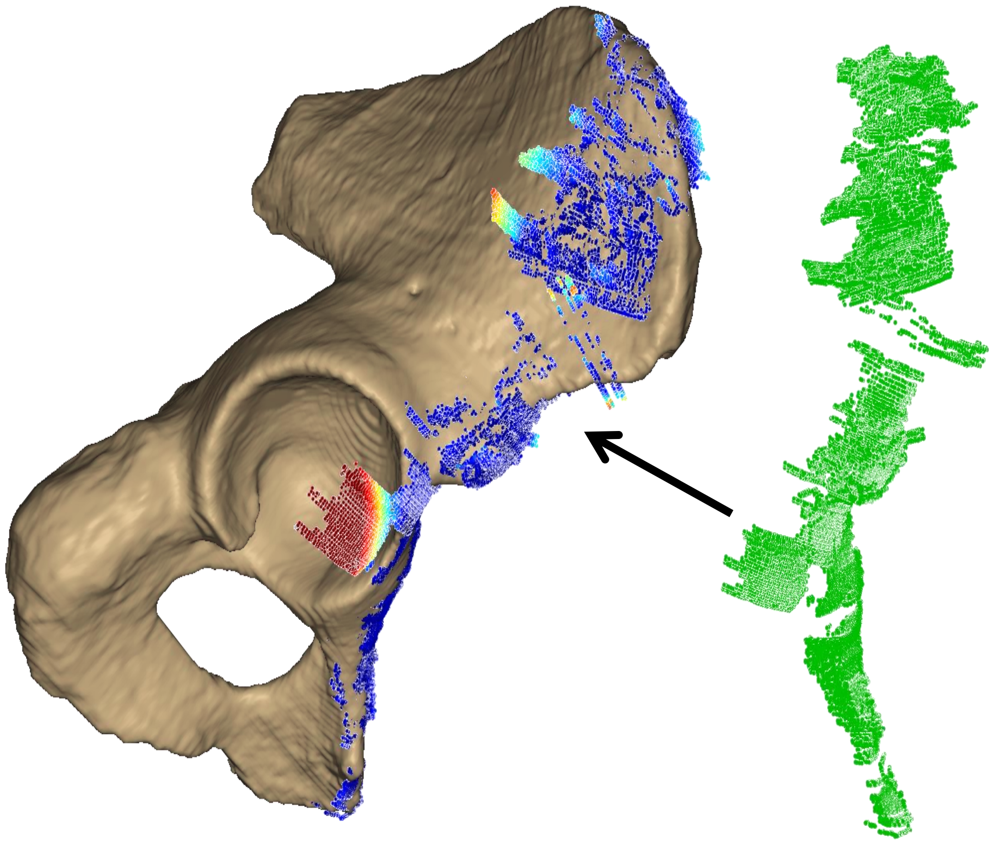

Point Cloud to Surface Registration

Optimization problem

Minimize the distance between each point and the closest point on the surface

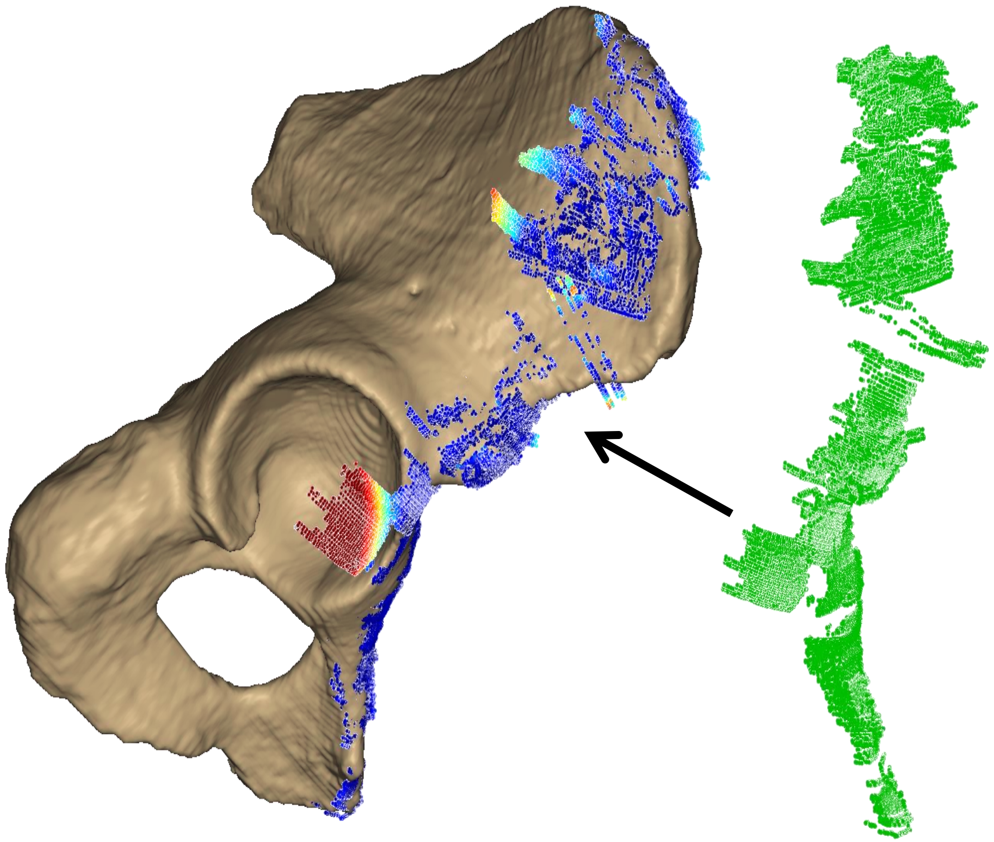

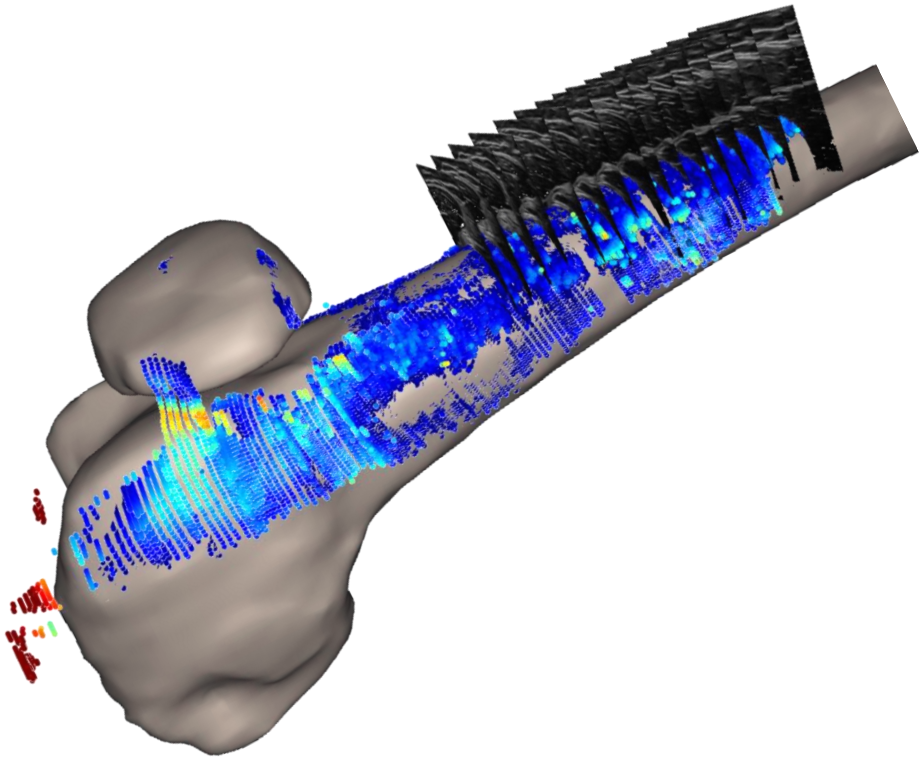

Fusion with pre-operative image

Salehi & Prevost et al.

Precise Ultrasound Bone Registration with Learning-Based Segmentation and Speed of Sound Calibration

MICCAI 2017

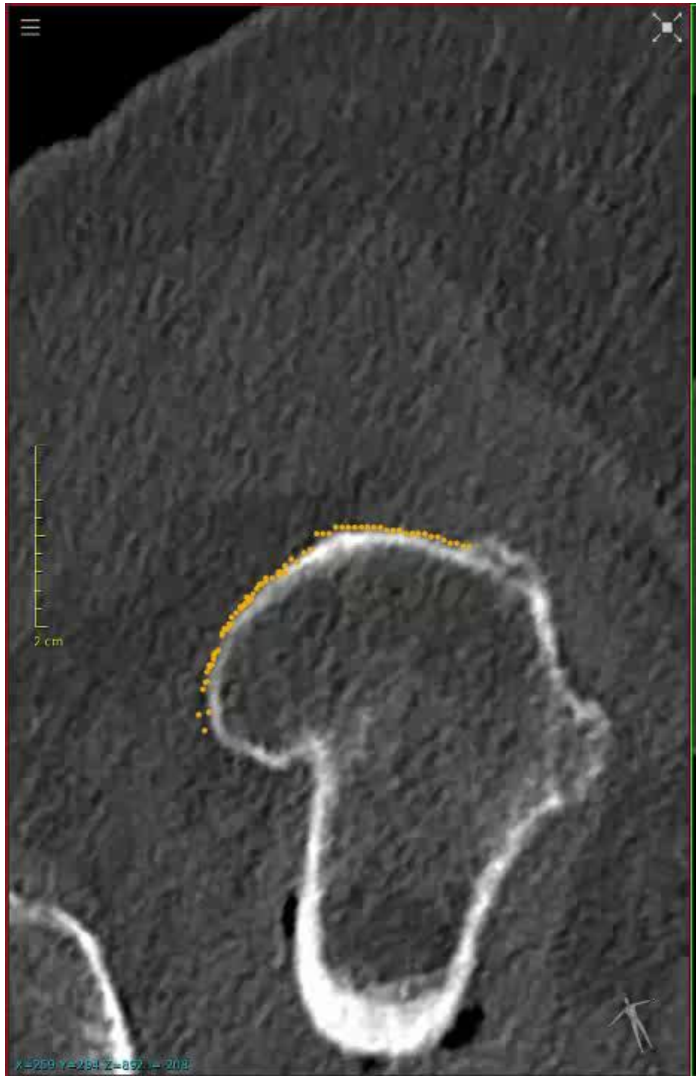

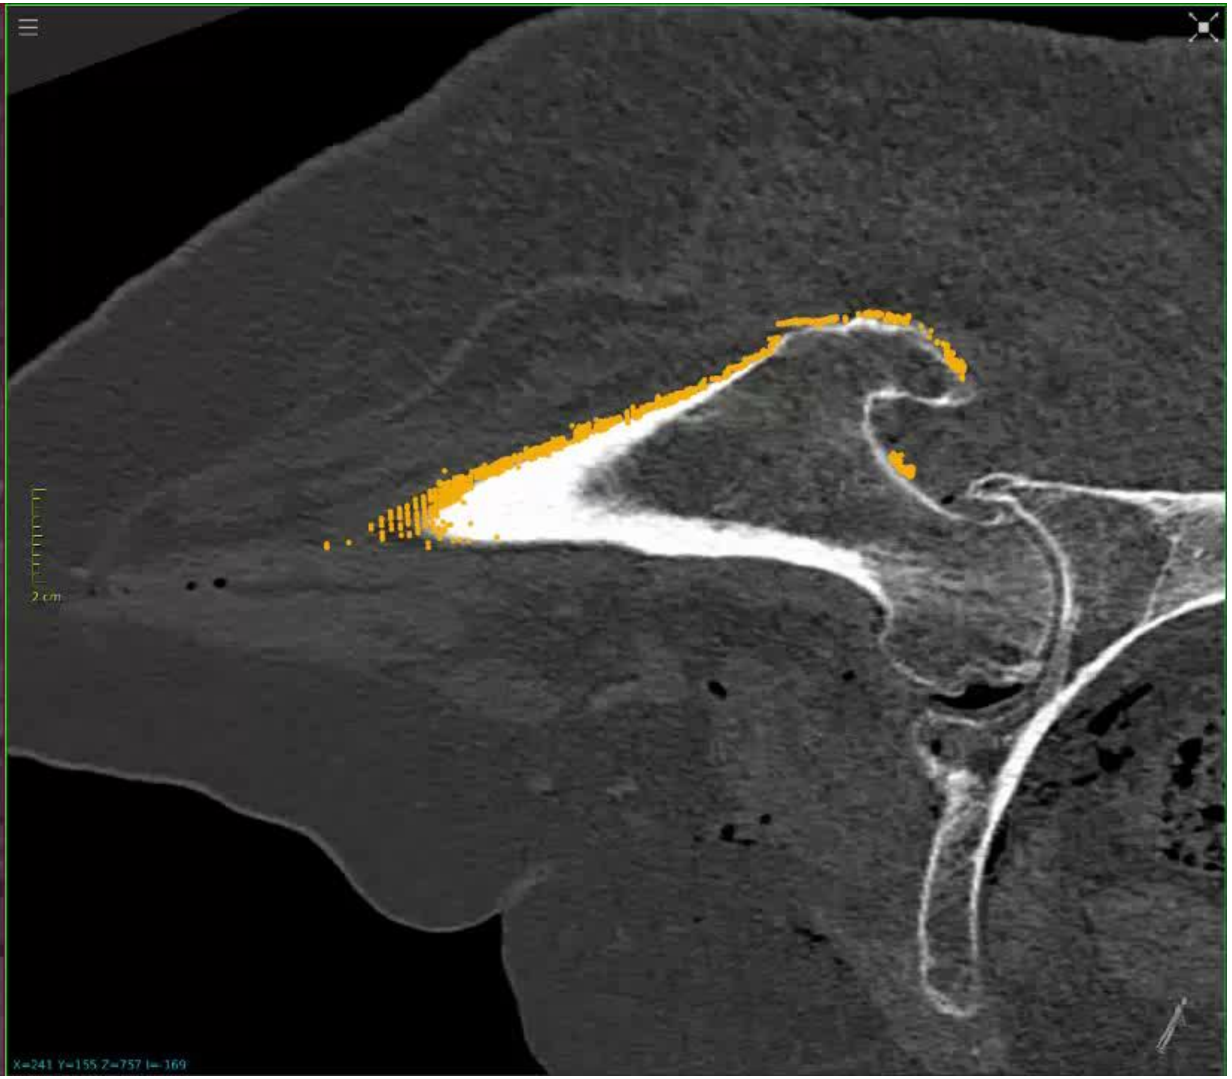

Extension to multiple bones

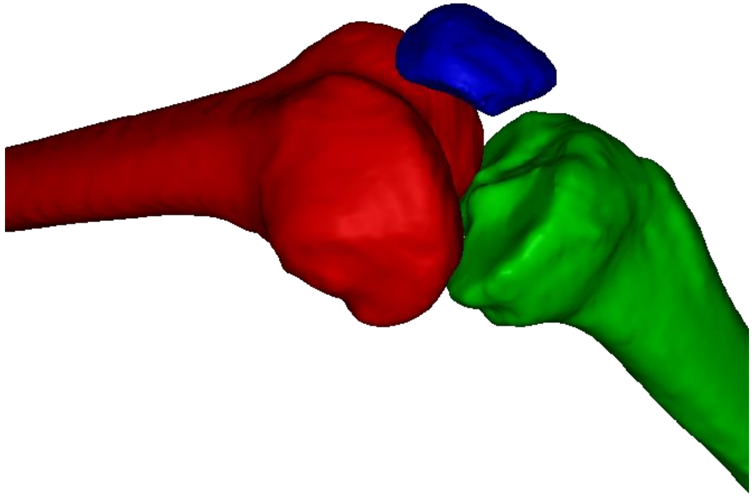

Extension to multiple bones

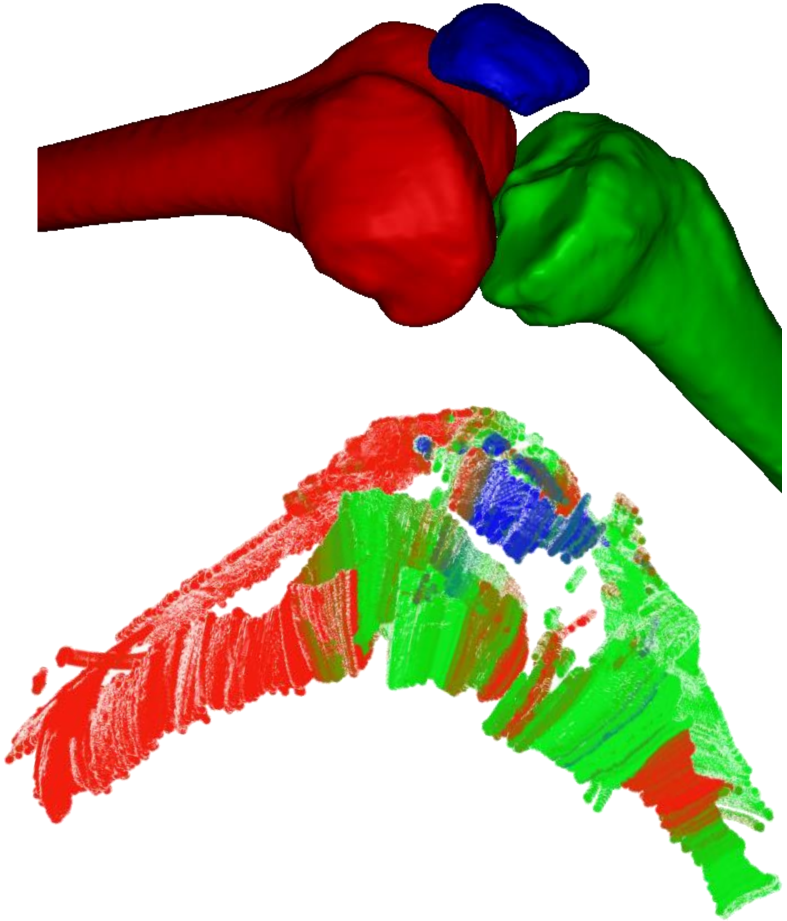

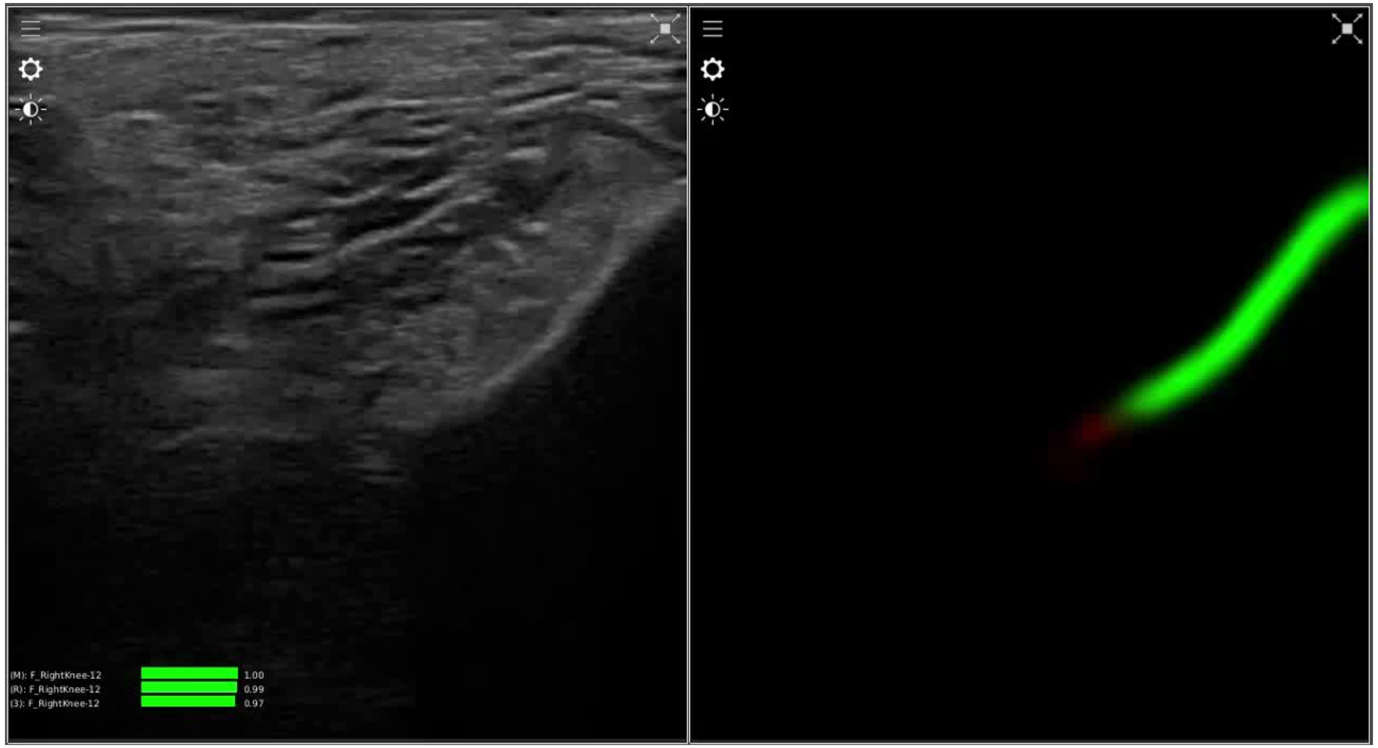

Train a neural network on different bones separately by encoding them as multiple channels

Extension to multiple bones

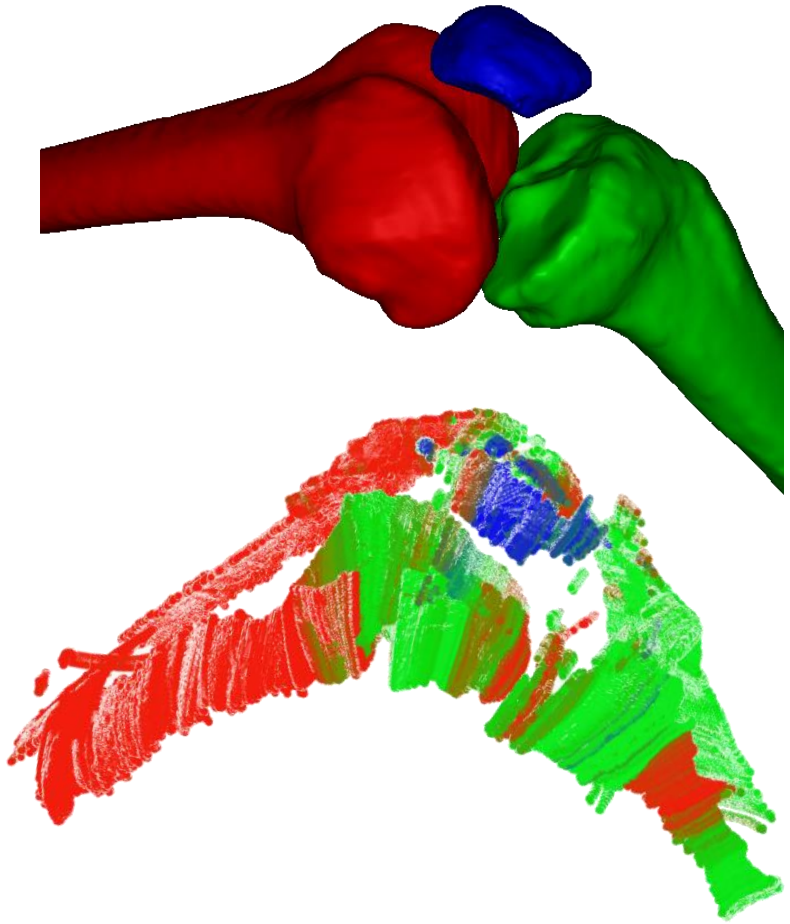

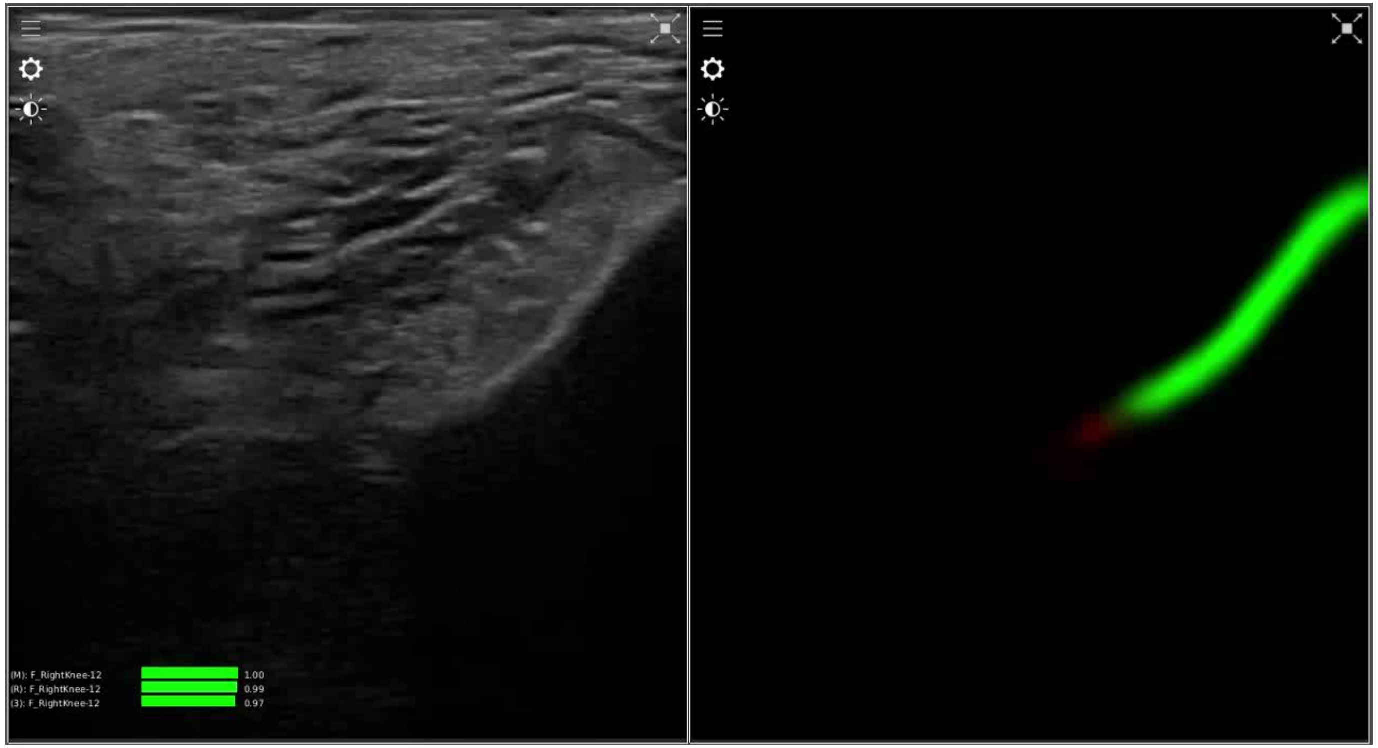

Train a neural network on different bones separately
by encoding them as multiple channels

Optimize the system accuracy by leveraging AI

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized

Optimize the system accuracy by leveraging AI

The who

✓ Ac

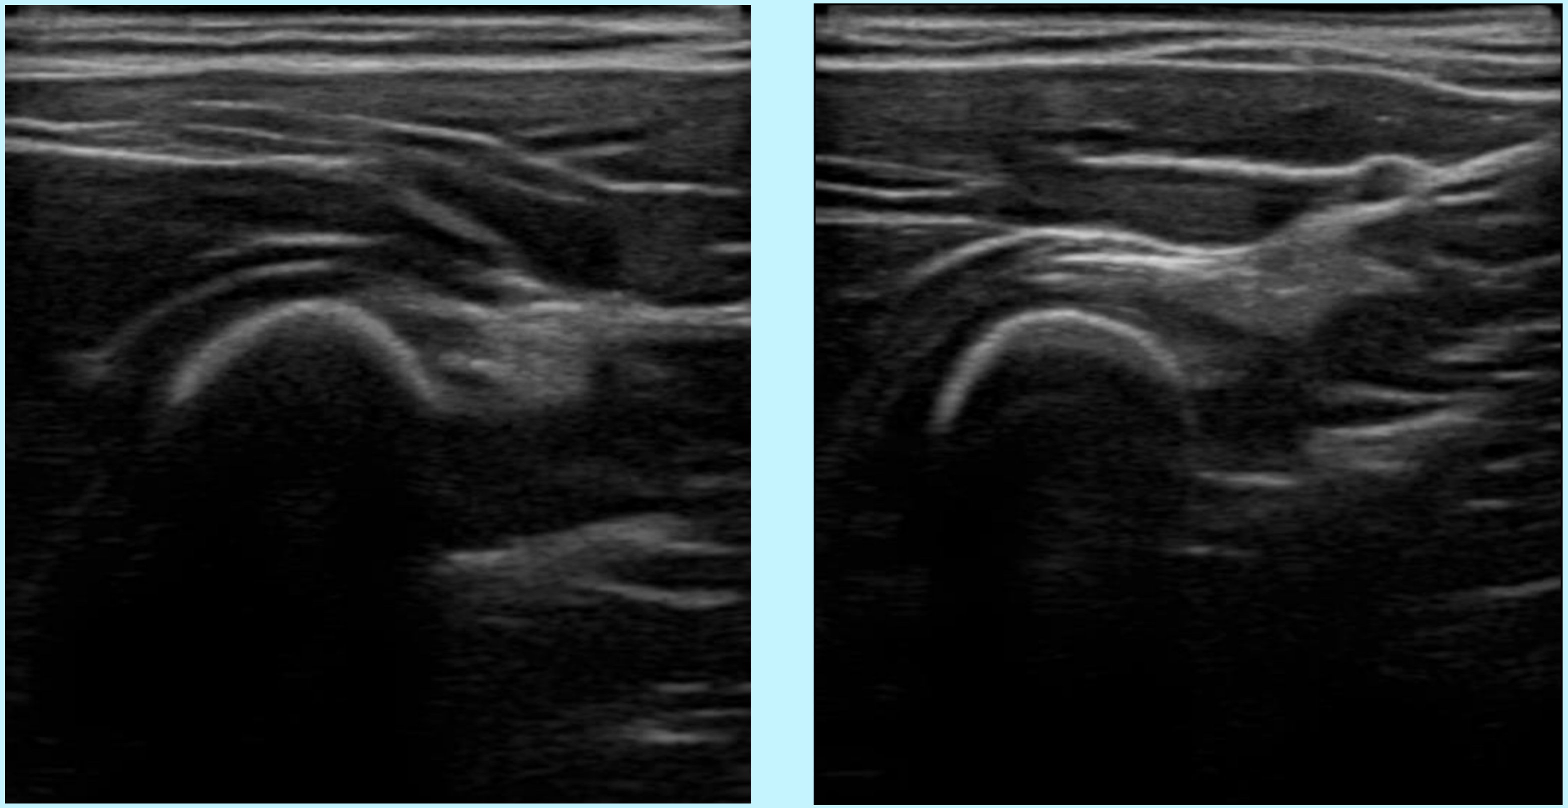

Bone surface can be more or less fuzzy

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally

Optimize the system accuracy by leveraging AI

Tracking camera

Temporal delay between images and position information

How to synchronize them ?

Ultrasound system

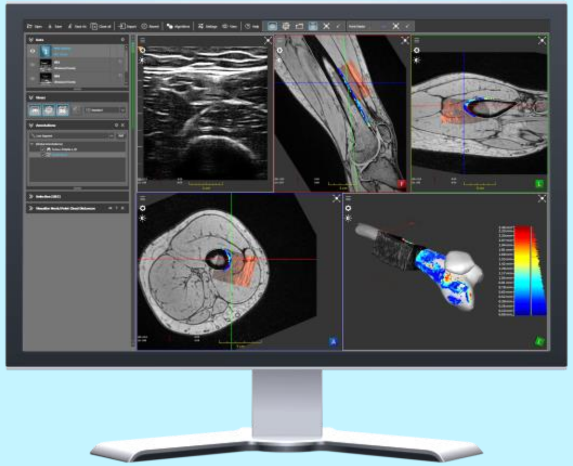

Position

Images

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally
- ✓ Speed of sound must be compensated

Optimize the system accuracy by leveraging AI

US systems assume a constant speed of sound
However, sound travels at different speeds in fat and muscle

US System Assumption
1540 m/s
Fat
1470 m/s
Muscle
1620 m/s

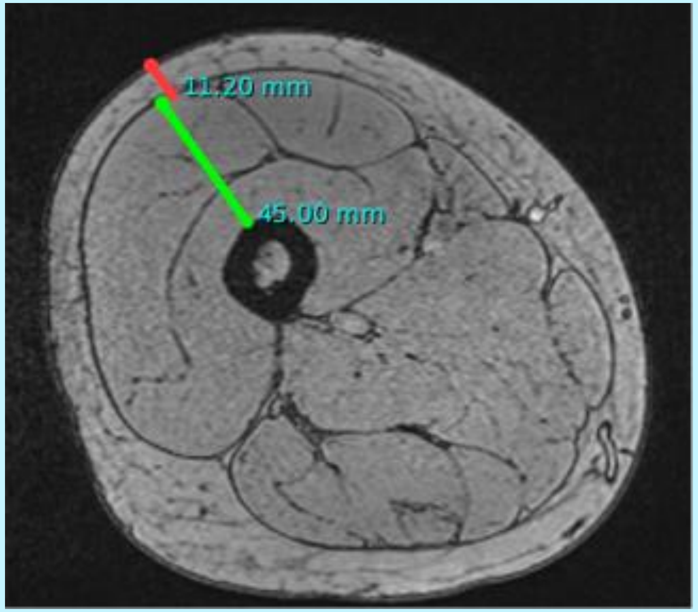

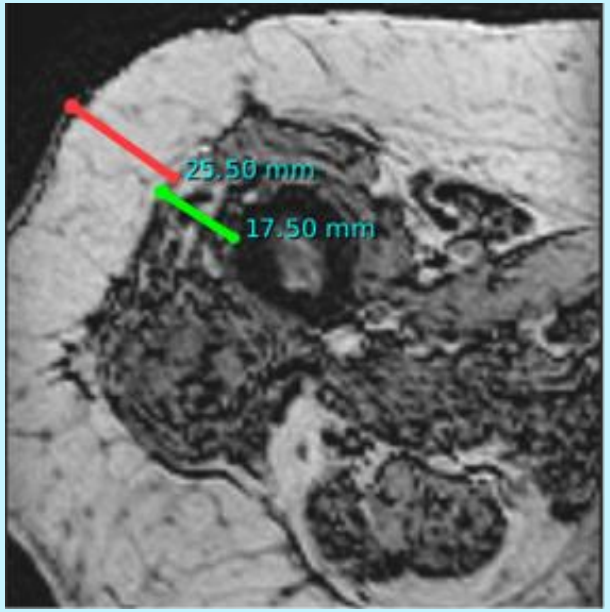

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally
- ✓ Speed of sound must be compensated

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally
- ✓ Speed of sound must be compensated

Such processes are usually tedious and complex

Optimize the system accuracy by leveraging AI

The whole system needs to be precisely calibrated

- ✓ Acquisition parameters must be optimized
- ✓ System must be calibrated geometrically and temporally
- ✓ Speed of sound must be compensated

Such processes are usually tedious and complex

...but we can leverage our real-time algorithms to solve them!

1) Parameter Tuning - Auto-Focus for Cameras

Cameras can automatically find the region of interest in an image and optimize the acquisition parameters

1) Parameter Tuning - Auto-Focus for Cameras

Cameras can automatically find the region of interest in an image
and optimize the acquisition parameters

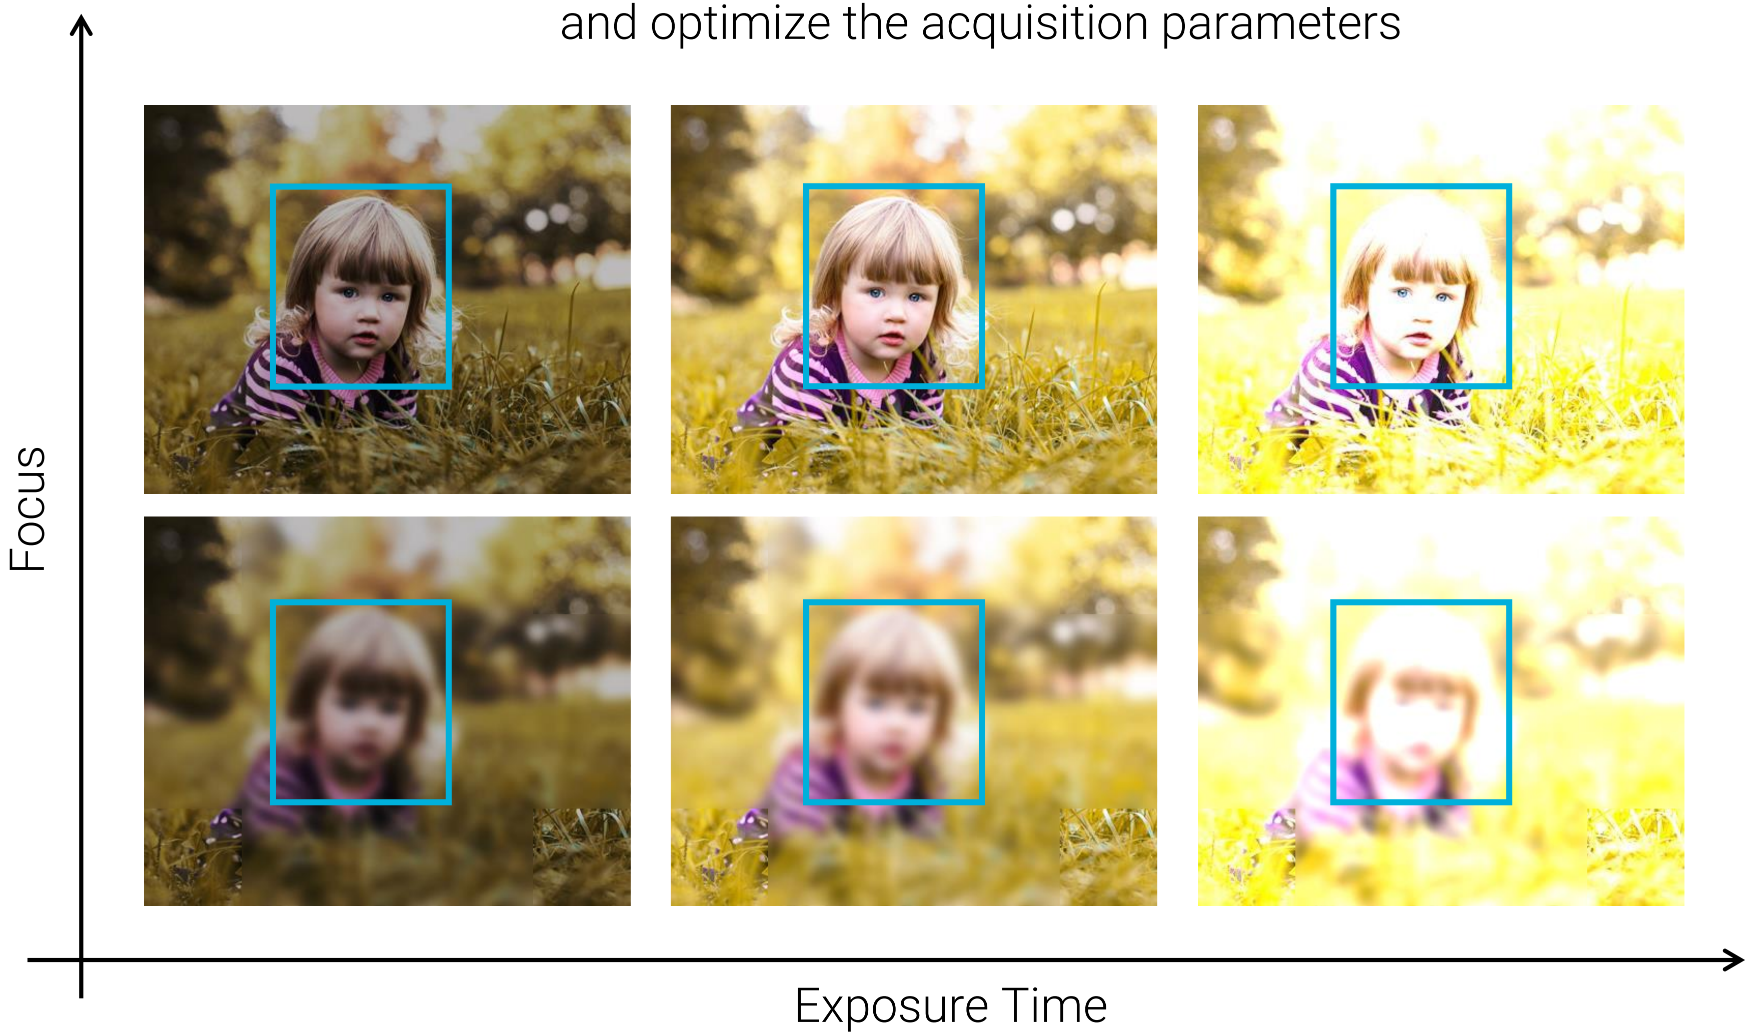

1) Parameter Tuning - Auto-Focus for Cameras

Cameras can automatically find the region of interest in an image and optimize the acquisition parameters

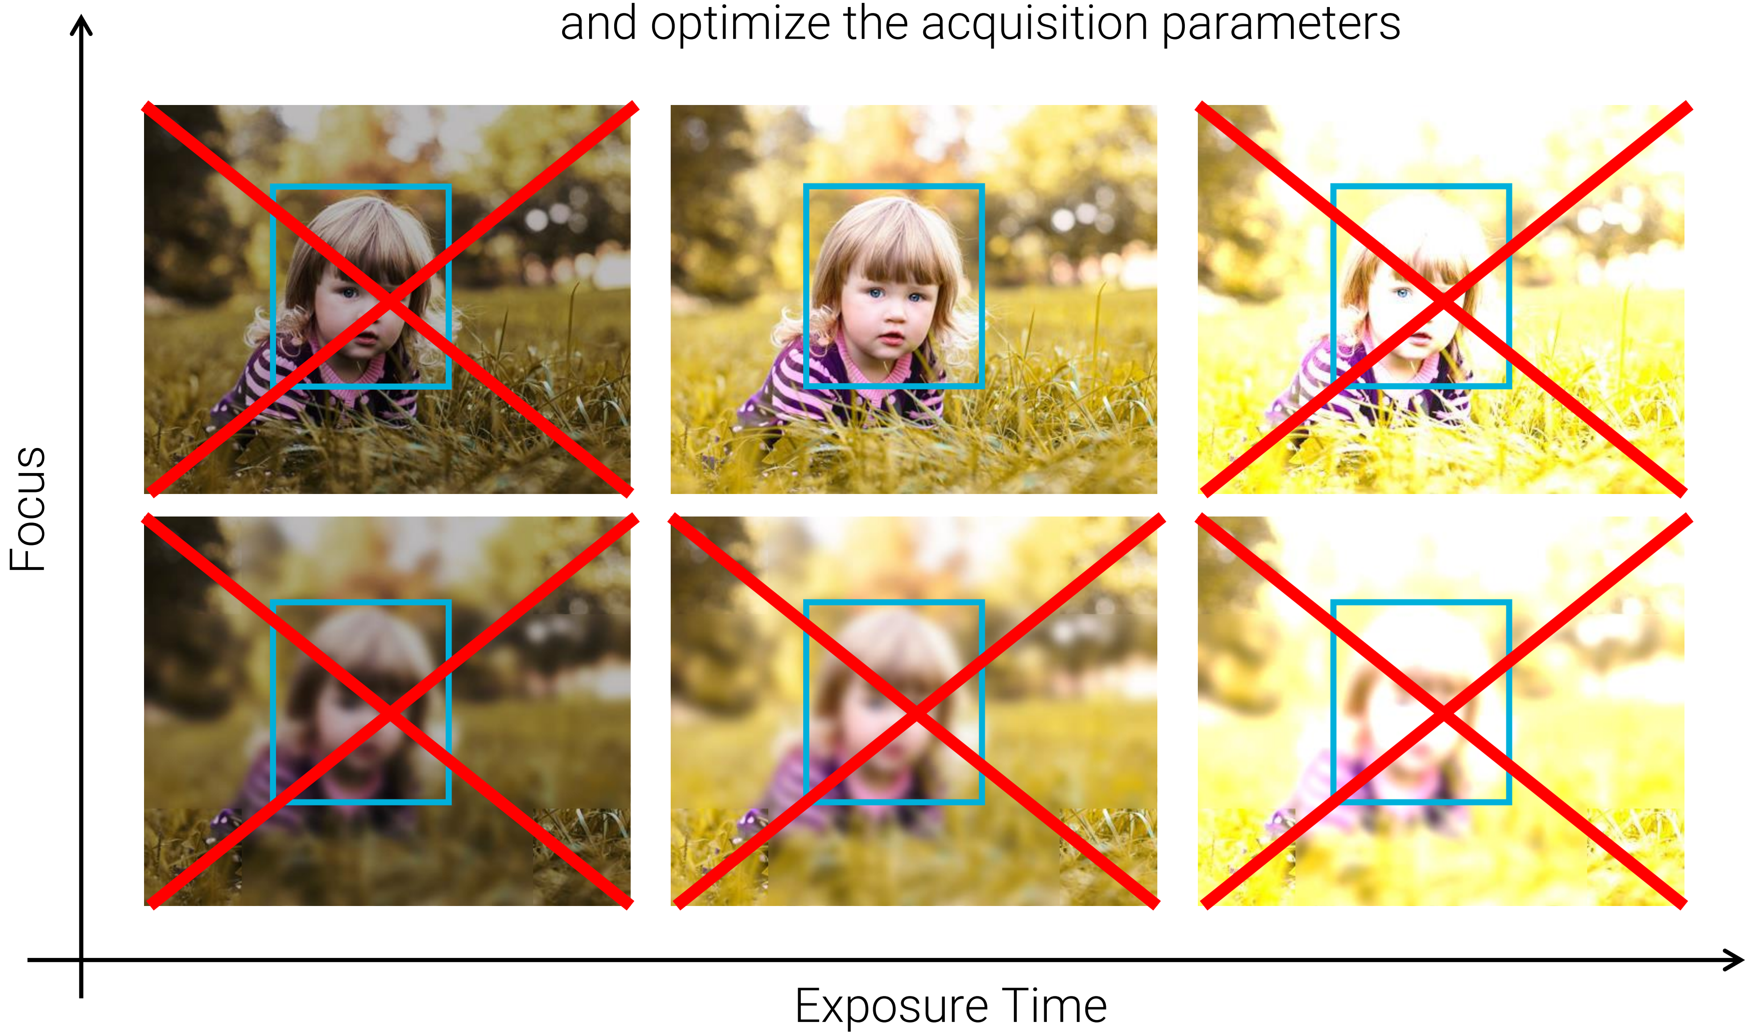

1) Parameter Tuning - Auto-Focus for Ultrasound!

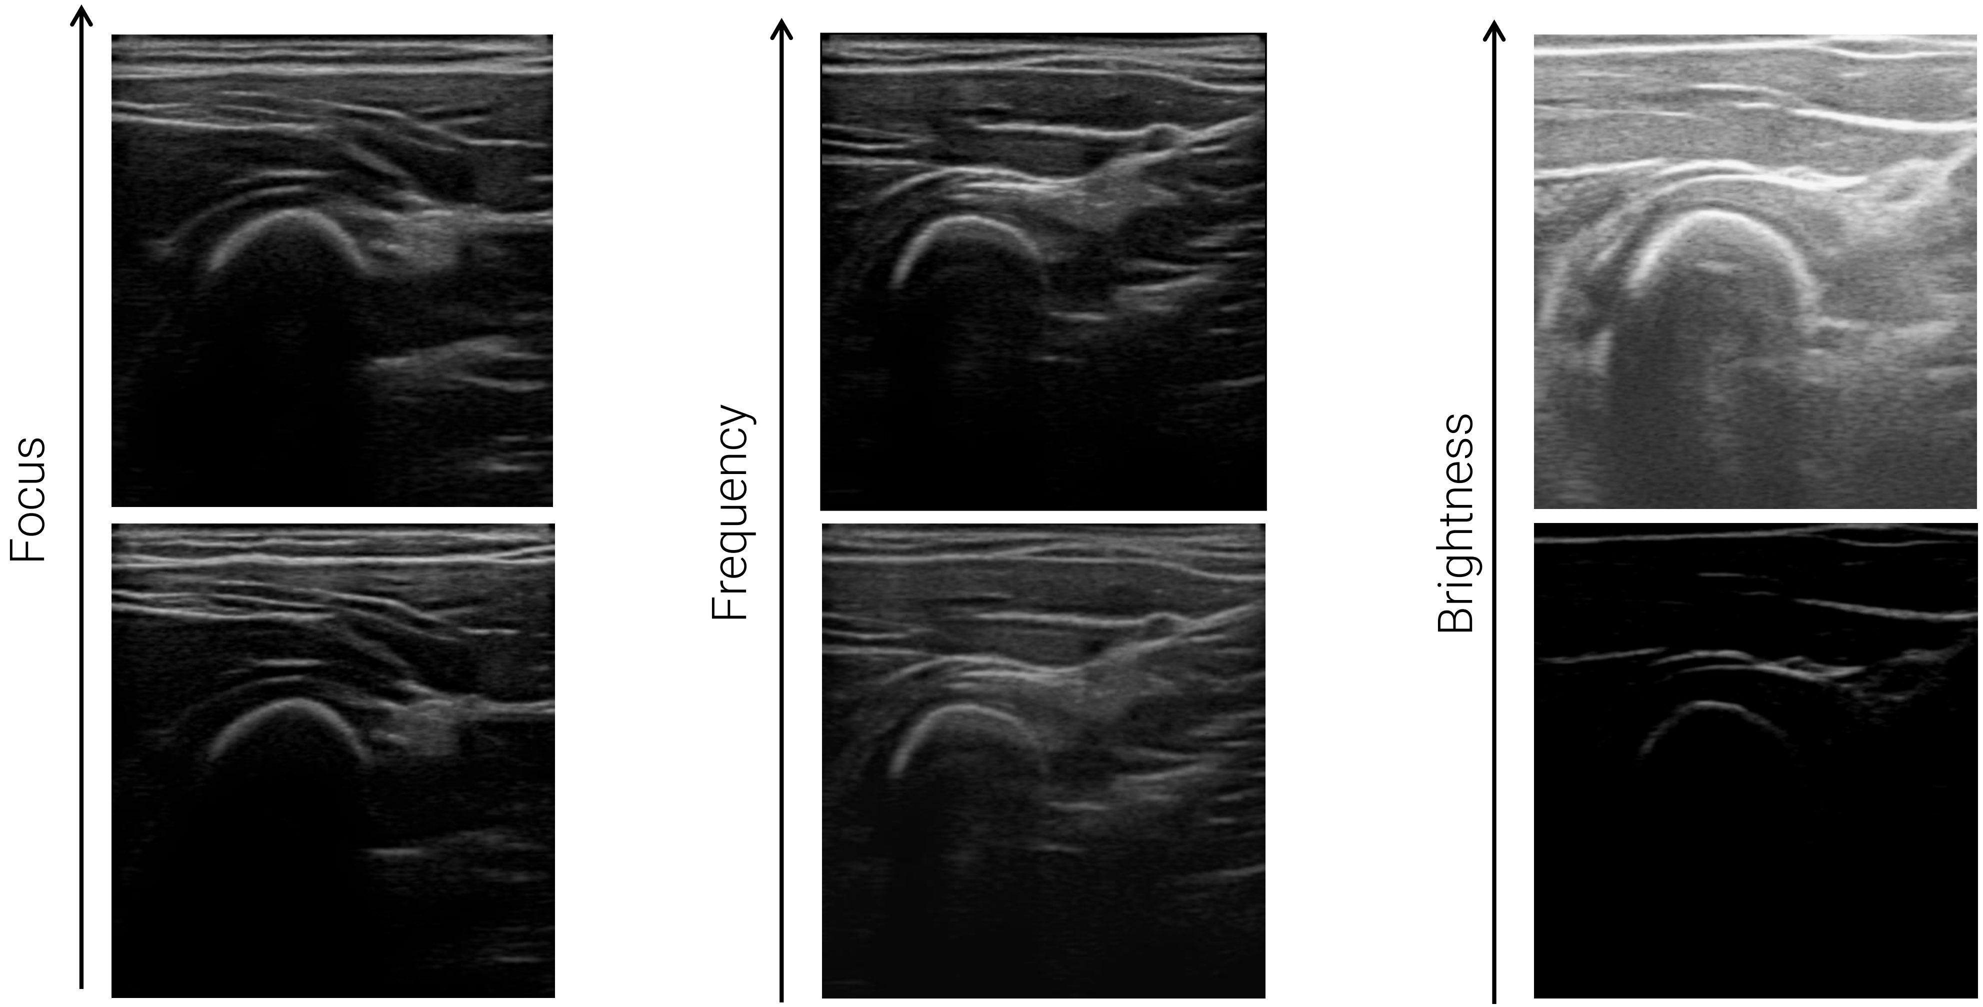

1) Parameter Tuning - Auto-Focus for Ultrasound!

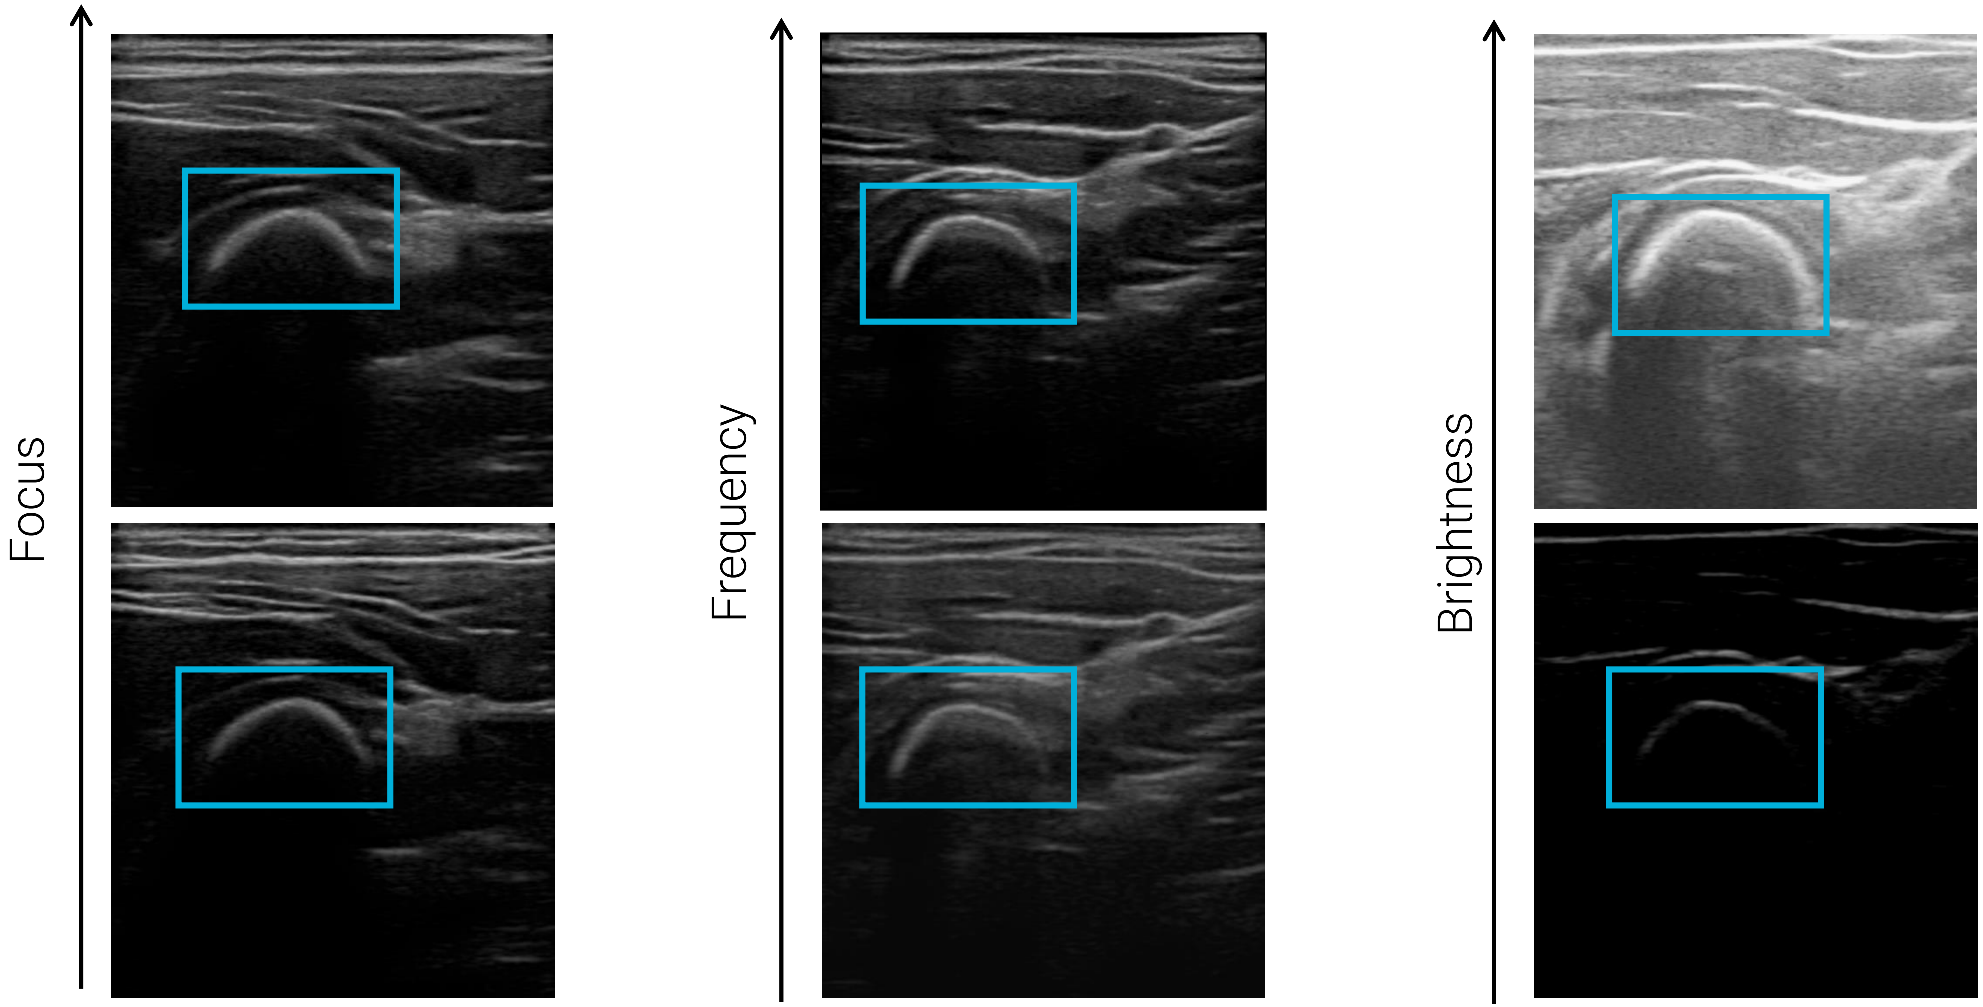

1) Parameter Tuning - Auto-Focus for Ultrasound!

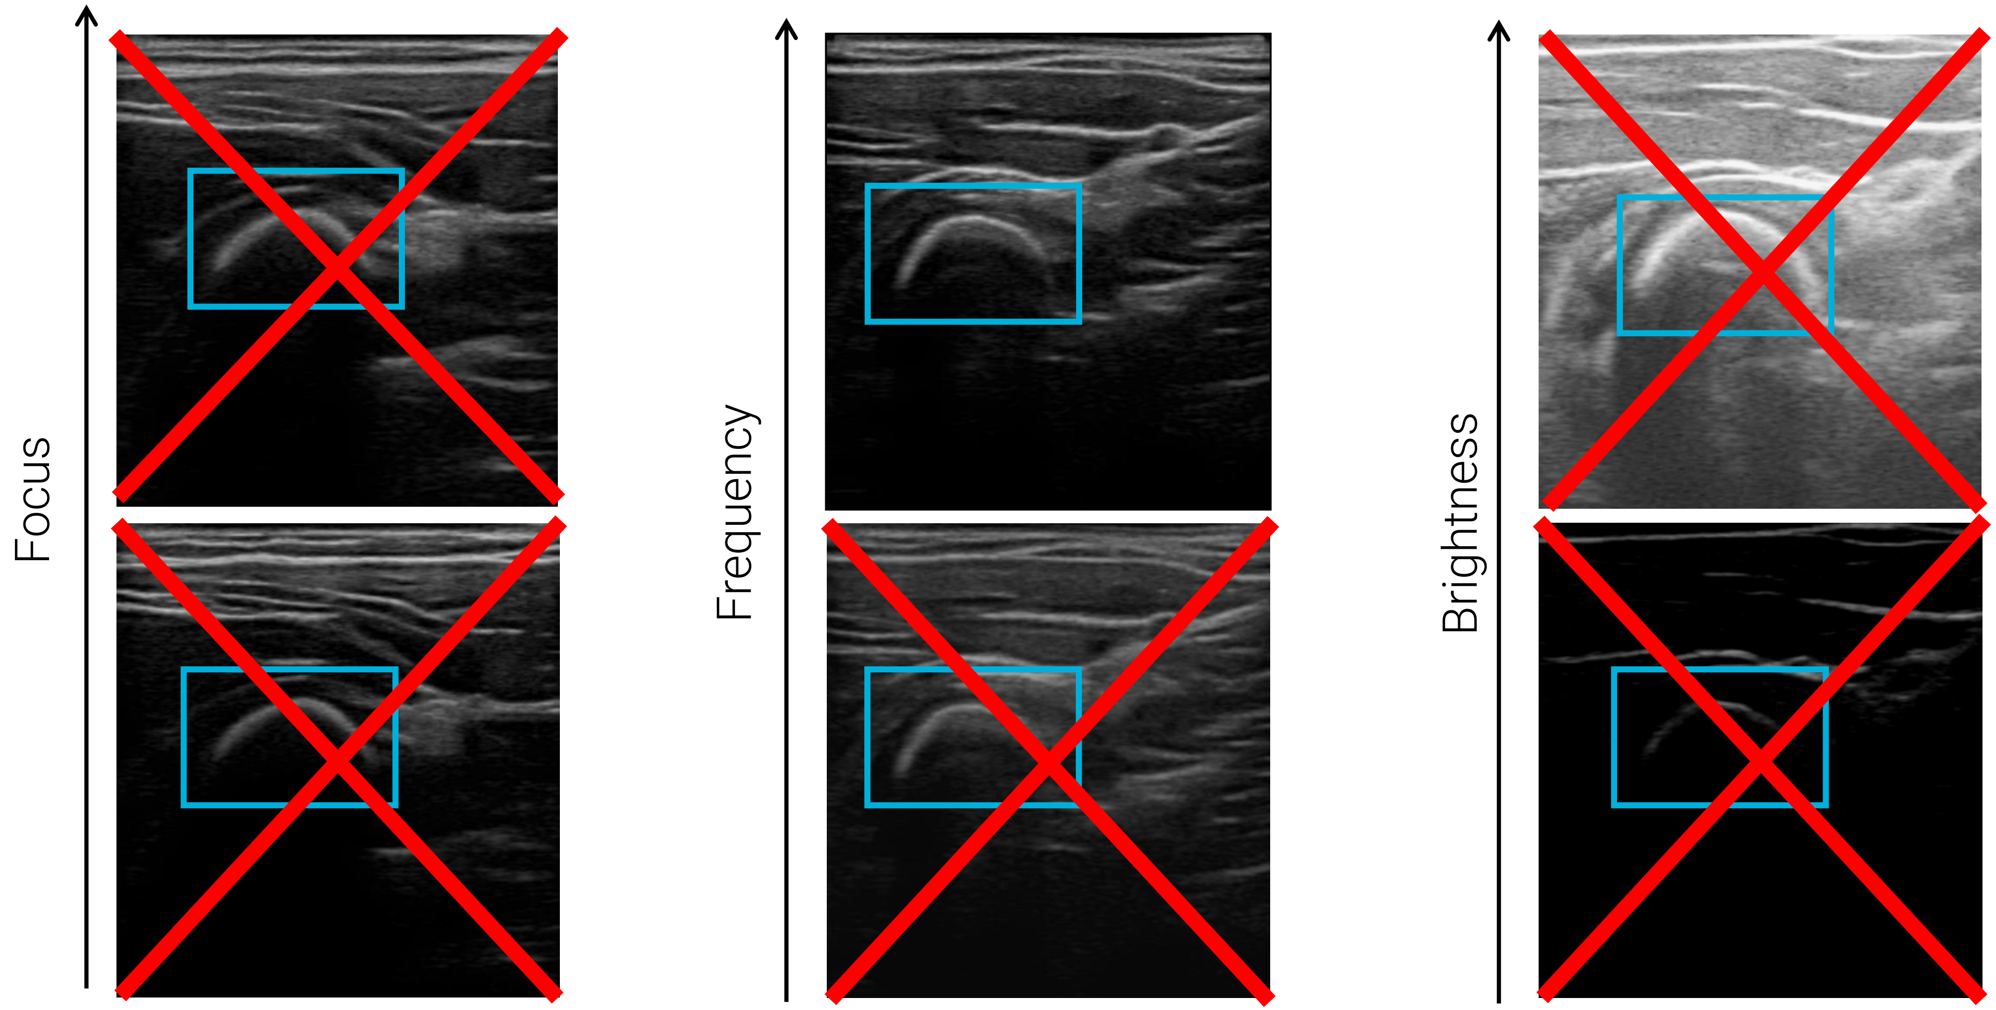

1) Automatic Acquisition Parameter Tuning

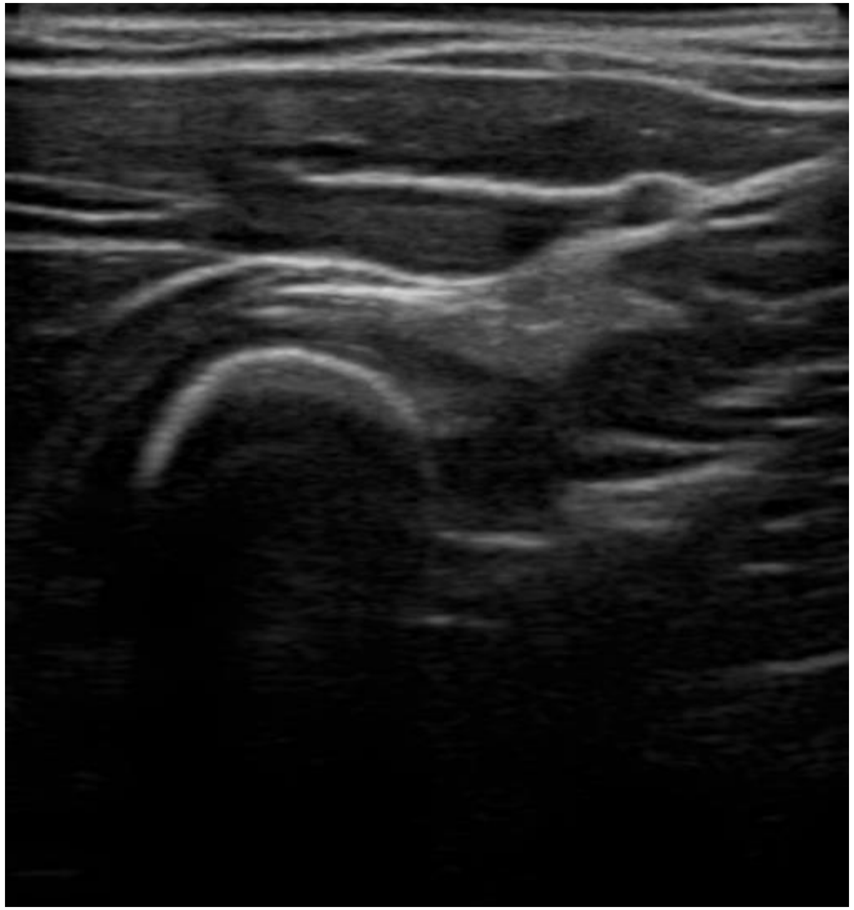

1) Automatic Acquisition Parameter Tuning

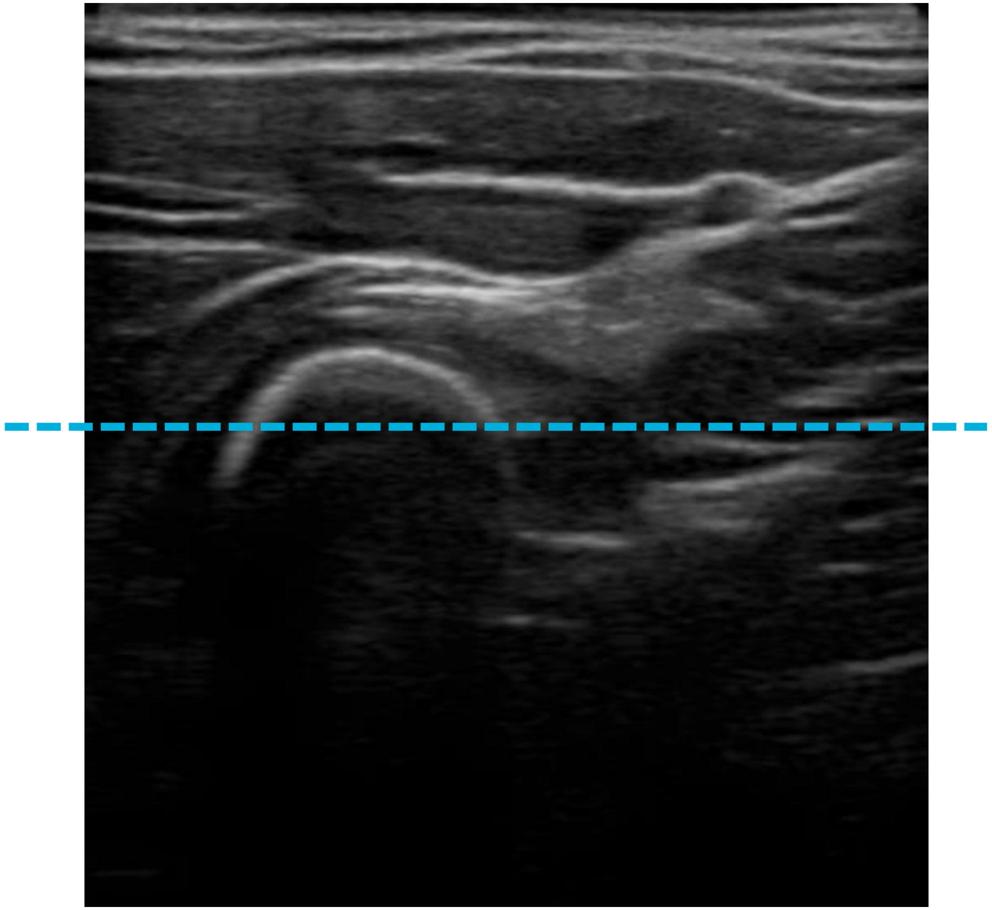

- Focus is equal to the depth of the bone

1) Automatic Acquisition Parameter Tuning

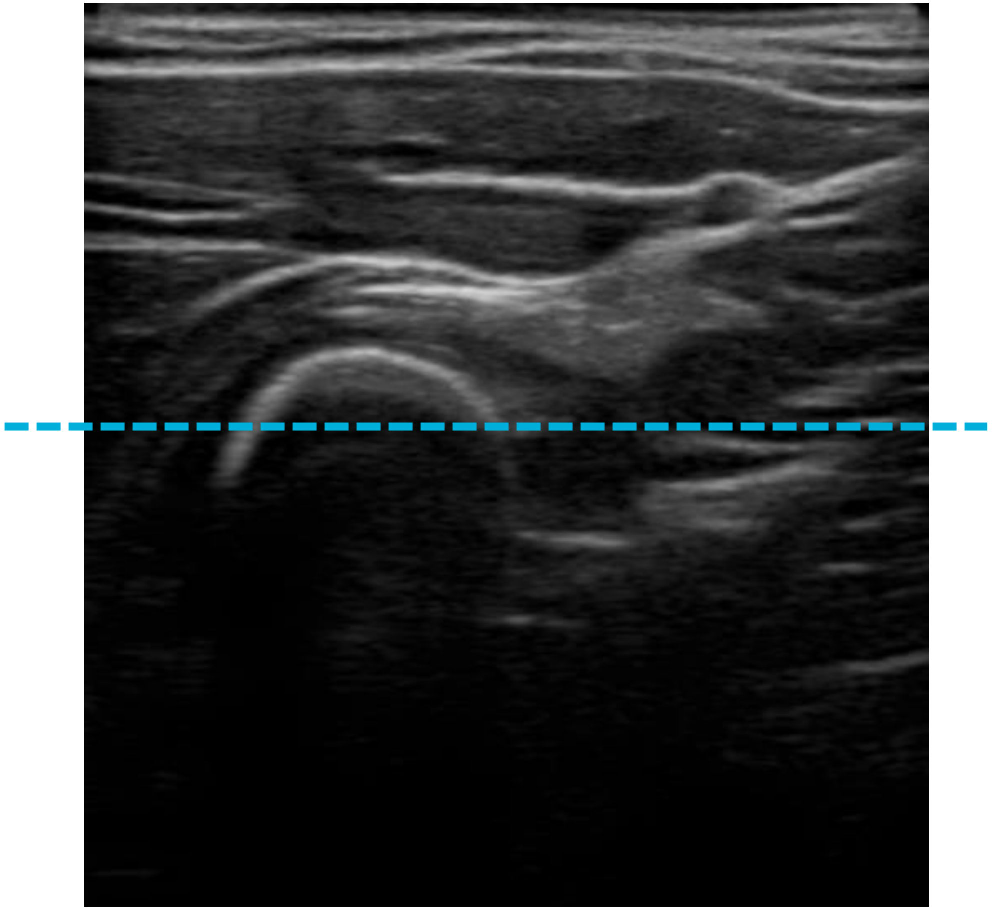

- Focus is equal to the depth of the bone

- Frequency also depends on the depth of the bone (high frequencies do not travel deep enough)

1) Automatic Acquisition Parameter Tuning

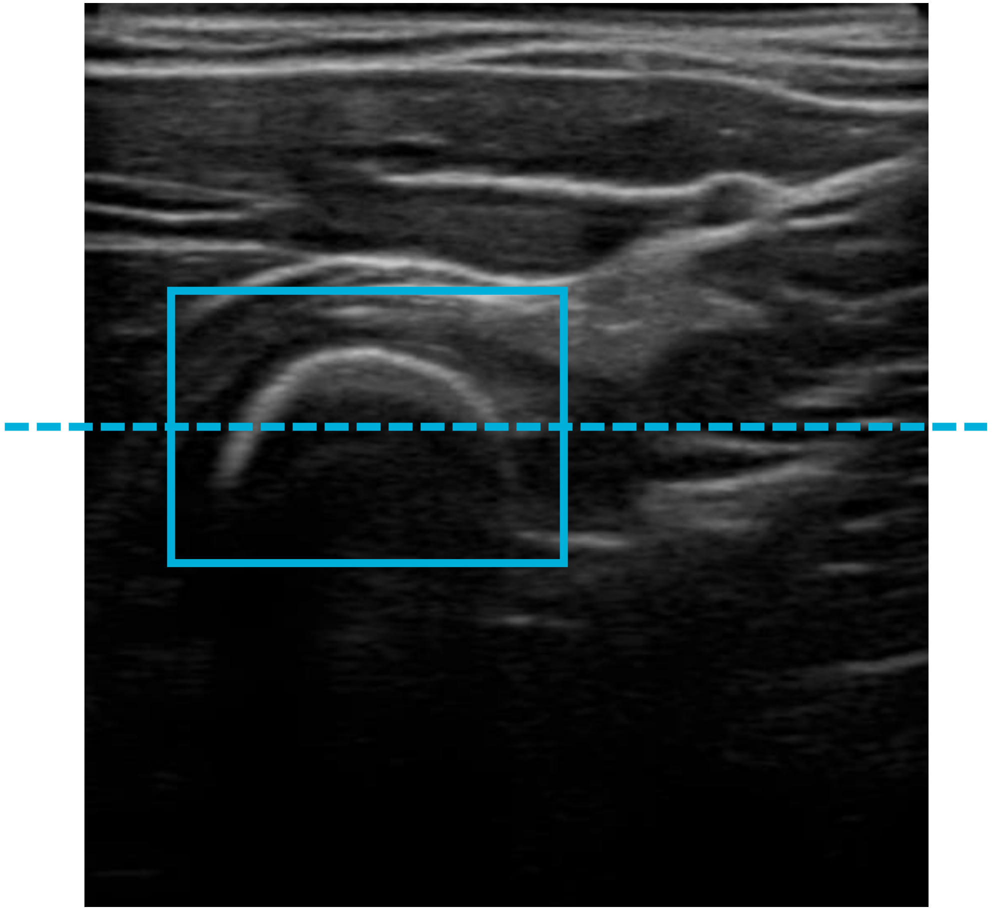

- Focus is equal to the depth of the bone
- Frequency also depends on the depth of the bone (high frequencies do not travel deep enough)
- Brightness can be adjusted by computing intensity statistics

Live Demo
Auto-Focus

in partnership with

Cephasonics Ultrasound Solutions

2) Calibrations

Speed of sound correction

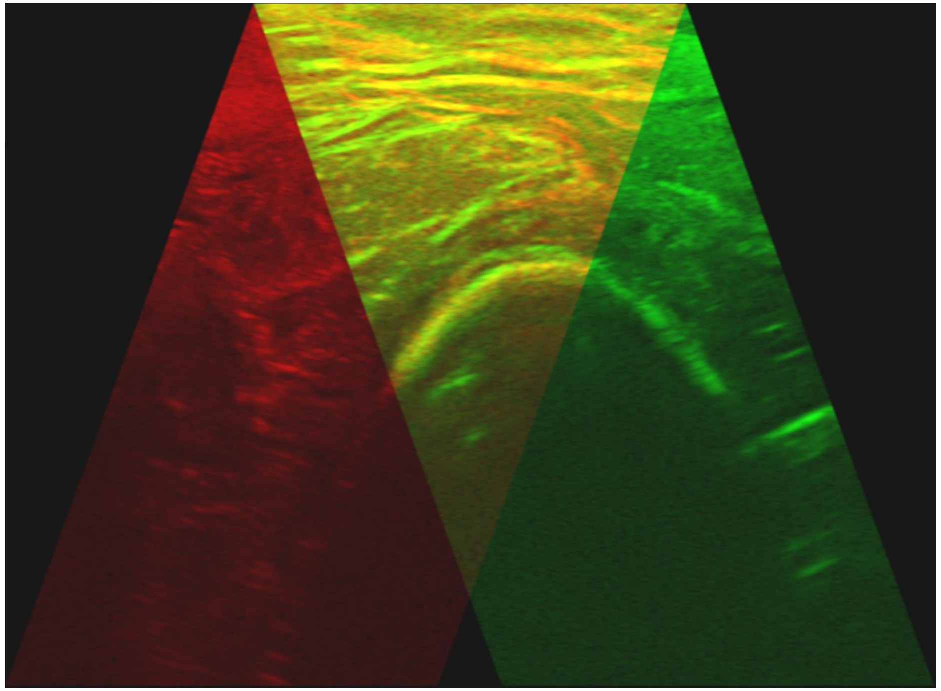

Temporal calibration

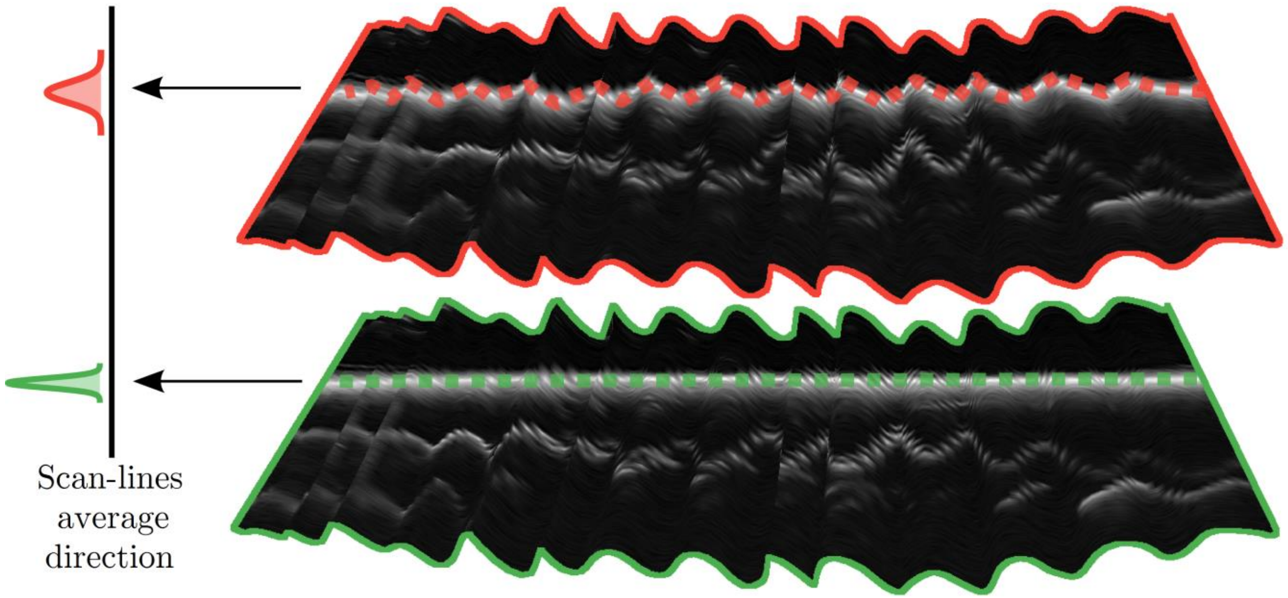

Salehi & Prevost et al.
Precise Ultrasound Bone Registration with Learning-Based Segmentation and Speed of Sound Calibration
MICCAI 2017

PART 3
NEURO SURGERY

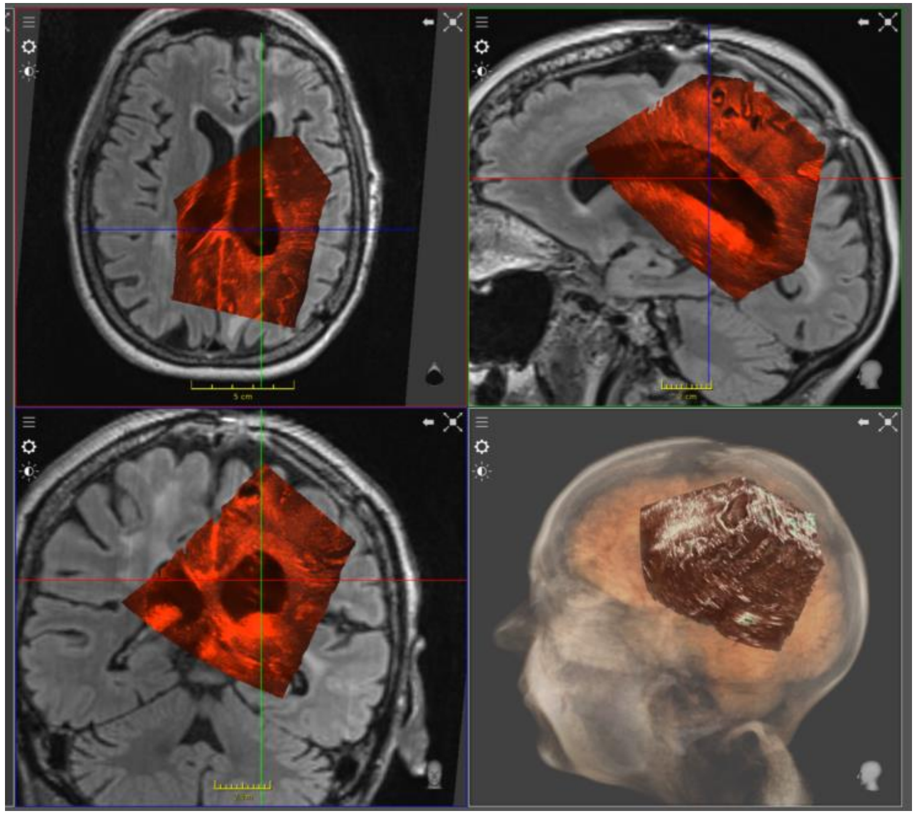

From planning to brain surgery

- Brain surgery usually planned on pre-operative MRI
  Where is the tumor? How big is it?

From planning to brain surgery

- Brain surgery usually planned on pre-operative MRI
  Where is the tumor? How big is it?

- In the OR, very difficult to follow a surgical plan

- **Brain shift:** When the skull is opened, gravity causes the brain to collapse

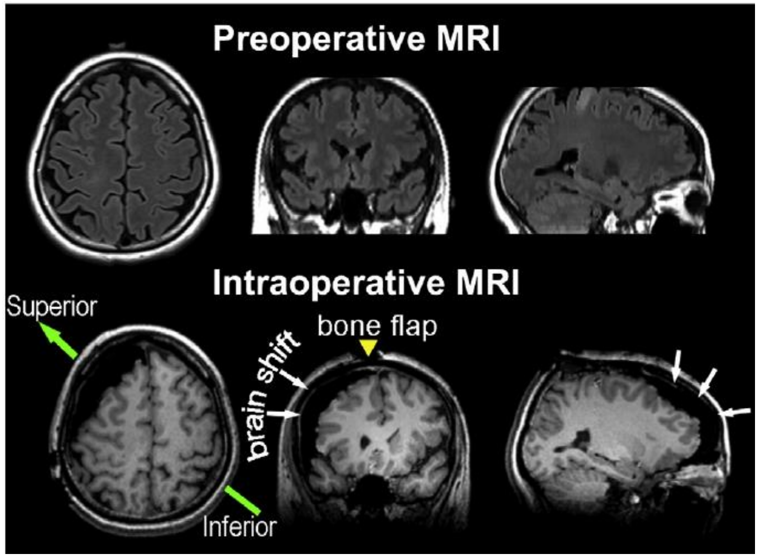

Lu, Jun-Feng, et al. *NeuroImage: Clinical* 2 (2013): 132-142

From planning to brain surgery

- Brain surgery usually planned on pre-operative MRI
  Where is the tumor? How big is it?

- In the OR, very difficult to follow a surgical plan

- **Brain shift:** When the skull is opened, gravity causes the brain to collapse

- Idea: Acquire ultrasound during surgery
  Deformable registration to the MR image
  → Planning can be used

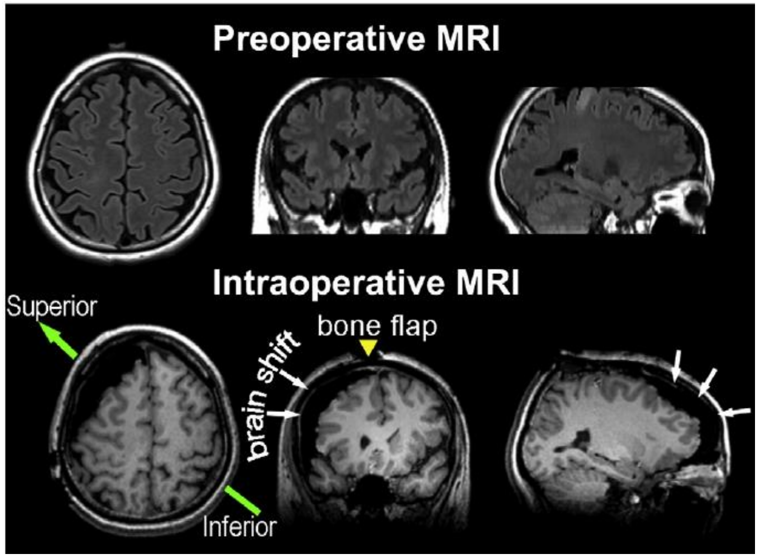

Lu, Jun-Feng, et al. *NeuroImage: Clinical* 2 (2013): 132-142

From planning to brain surgery

- Brain surgery usually planned on pre-operative MRI
  Where is the tumor? How big is it?

- In the OR, very difficult to follow a surgical plan

- **Brain shift:** When the skull is opened, gravity causes the brain to collapse

- Idea: Acquire ultrasound during surgery
  Deformable registration to the MR image
  → Planning can be used

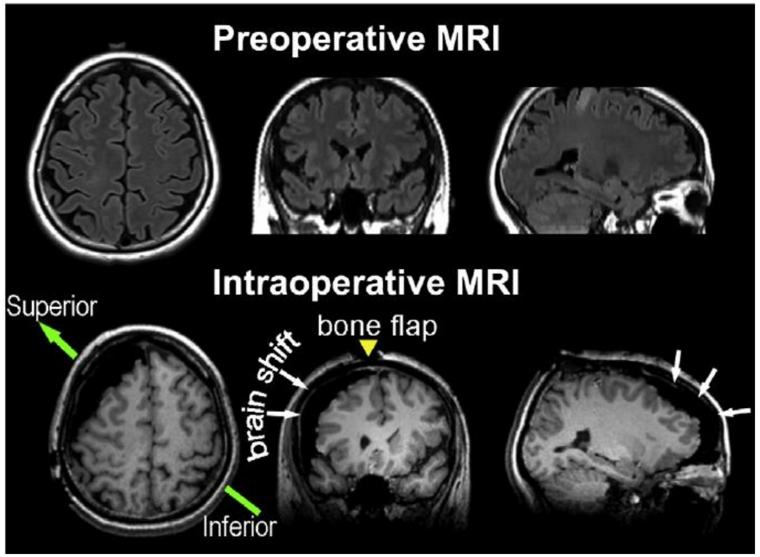

Lu, Jun-Feng, et al. *NeuroImage: Clinical* 2 (2013): 132-142

MRI to 3D Ultrasound Registration

MRI to 3D Ultrasound Registration

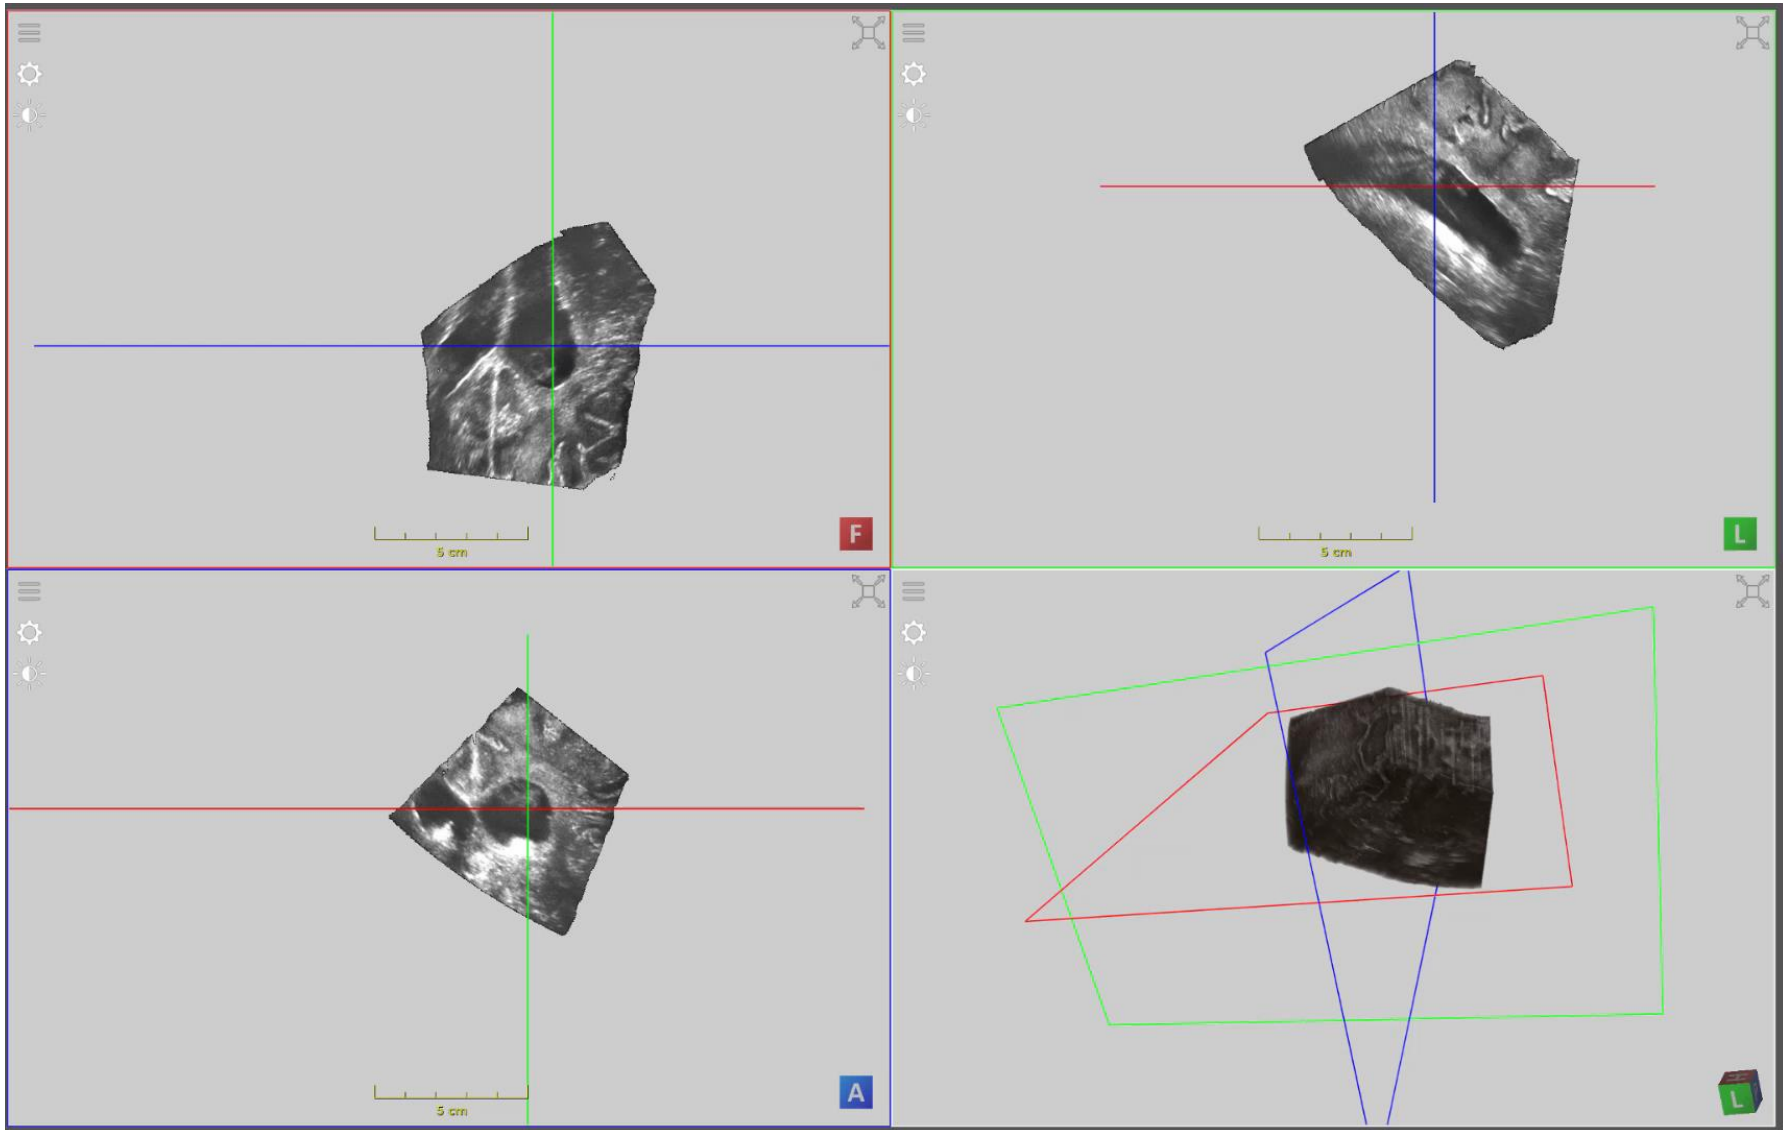

MICCAI CuRIOUS Challenge 2018

Correction of Brainshift with Intra-Operative Ultrasound
https://curious2018.grand-challenge.org

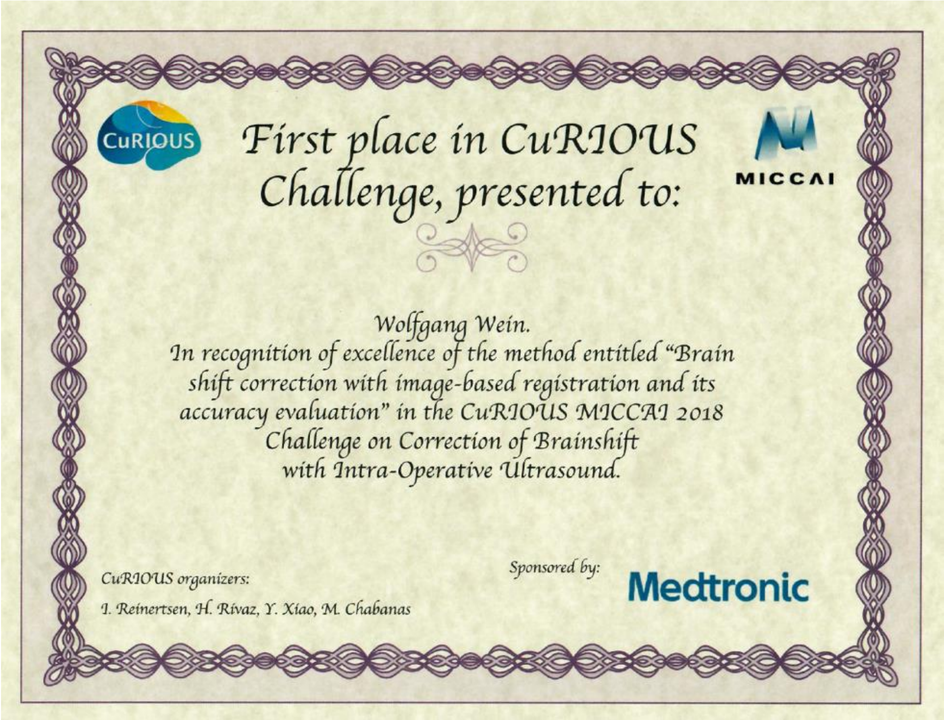

Not an AI method !

Similarity Measure: Instead of correlating US intensities with two channels of simulated information from CT as in [13], we use LC^2 to correlate US with both the MRI intensity values p and its spatial gradient magnitude $g = |\nabla p|$. The local LC^2 value is computed for each pixel $\mathbf{x_i}$ in each ultrasound image, considering a neighborhood $\Omega(\mathbf{x_i})$ of m pixels. For each patch of m pixels, the contribution of MRI intensity values p and gradient magnitudes g are unknown. Therefore, we define an intensity function $f(\mathbf{x_i})$ as a function of the transformed MRI intensities $p_i = p(T(\mathbf{x_i}))$ and gradients $g_i = g(T(\mathbf{x_i})) = |\nabla p_i|$ as:

$$f(\mathbf{x_i}) = \alpha p_i + \beta g_i + \gamma, \qquad (1)$$

where $y_i = \{\alpha, \beta, \gamma\}$ denotes the unknown parameters of the influence of the MRI intensities and gradients within $\Omega(\mathbf{x_i})$. They can be estimated by minimizing the difference of the intensity function and the ultrasound image intensity u_i:

$$\left\| \mathrm{M} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} - \begin{pmatrix} u_1 \\ \vdots \\ u_m \end{pmatrix} \right\|^2 \quad \text{where } \mathrm{M} = \begin{pmatrix} p_1 & g_1 & 1 \\ \vdots & \vdots & \vdots \\ p_m & g_m & 1 \end{pmatrix}, \qquad (2)$$

which can be solved using ordinary least squares with the pseudo-inverse of M. This results in a parameter triple y_i for each pixel $\mathbf{x_i}$, which is only depending on the neighborhood $\Omega(\mathbf{x_i})$ and therefore compensating for changing influences of tissue interfaces or organ-internal intensities. The local similarity is then:

$$S(u, M) = 1 - \frac{\sum_{\mathbf{x_i}} |u(\mathbf{x_i}) - My|^2}{\sum_{\mathbf{x_i}} Var(u(\mathbf{x_i}))} \qquad (3)$$

The overall similarity is the weighted sum of eq. 3 with the local variance of the US image. This suppresses regions without structural appearance, therefore allowing to cope with ultrasonic occlusions implicitly, without the need to simulate them.

Wein et al.
Global Registration of Ultrasound to MRI Using the LC2 Metric for Enabling Neurosurgical Guidance
MICCAI 2013

… but still computationally intensive→ GPU implementation

Top 3 methods were <u>not</u> based on machine learning

Part 4
Ultrasound For Vascular Imaging

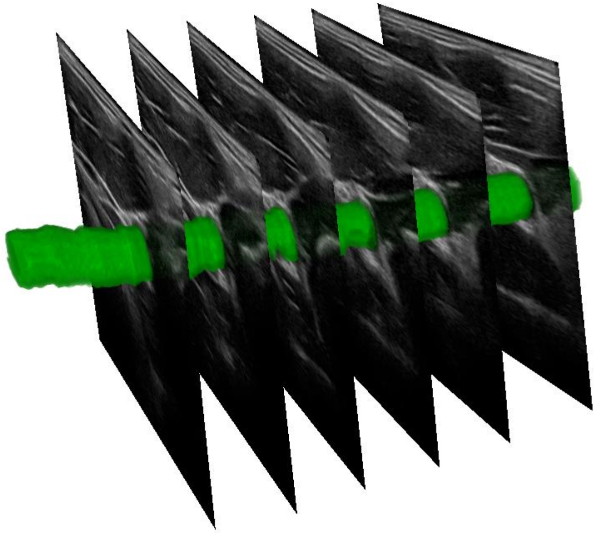

in partnership with

PIUR IMAGING

www.piurimaging.com

Vascular Imaging

- Visualization of blood vessels
- Multiple clinical applications, e.g.
  - Stenosis/Aneurysm Management and Surveillance
  - Fistula Planning and Monitoring
  - Vascular Mapping
- Typically performed with a CT or MR scanner after injection of contrast agents
  → Expensive, long, toxic
  → Not suited for screening or monitoring

MR

Vascular Imaging

- Visualization of blood vessels
- Multiple clinical applications, e.g.
  - Stenosis/Aneurysm Management and Surveillance
  - Fistula Planning and Monitoring
  - Vascular Mapping
- Typically performed with a CT or MR scanner after injection of contrast agents
  → Expensive, long, toxic
  → Not suited for screening or monitoring

MR

US

source: piurimaging.com

From 2D to 3D US – without External Hardware

Existing Hardware Solutions

Motorized Transducer
“wobbler”

Limited field of view
Temporal artifacts

Matrix Array
“3d probe”

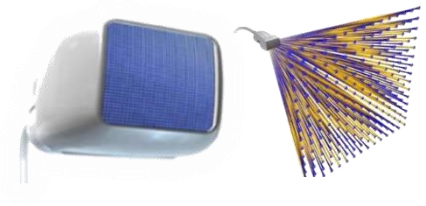

Philips xMatrix

Limited field of view
Decreased image quality

Tracking
(optical/EM)

Expensive
Not portable

Our Goal

Image-Based
Reconstruction

No hardware

From 2D to 3D US – without External Hardware

Existing Hardware Solutions

Motorized Transducer "wobbler"	Matrix Array "3d probe"	Tracking (optical/EM)
	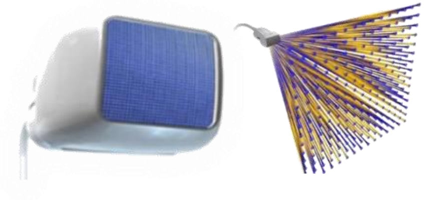	
Limited field of view Temporal artifacts	Limited field of view Decreased image quality	Expensive Not portable

---

Our Goal

Image-Based Reconstruction

No hardware

From 2D to 3D US – without External Hardware

Existing Hardware Solutions

Motorized Transducer "wobbler"	Matrix Array "3d probe"	Tracking (optical/EM)
	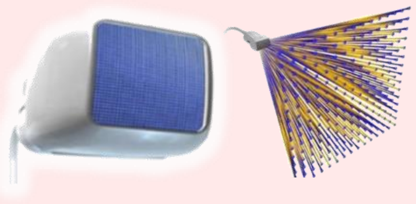 Philips xMatrix	
Limited field of view Temporal artifacts	Limited field of view Decreased image quality	Expensive Not portable

Our Goal

Image-Based Reconstruction

No hardware

From 2D to 3D US – without External Hardware

Existing Hardware Solutions

Motorized Transducer "wobbler"

Limited field of view
Temporal artifacts

Matrix Array "3d probe"

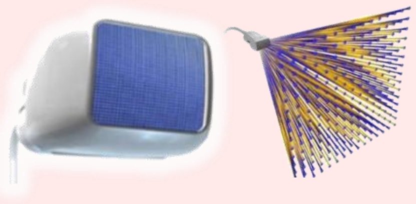

Limited field of view
Decreased image quality

Tracking (optical/EM)

Expensive
Not portable

Our Goal

Image-Based Reconstruction

No hardware

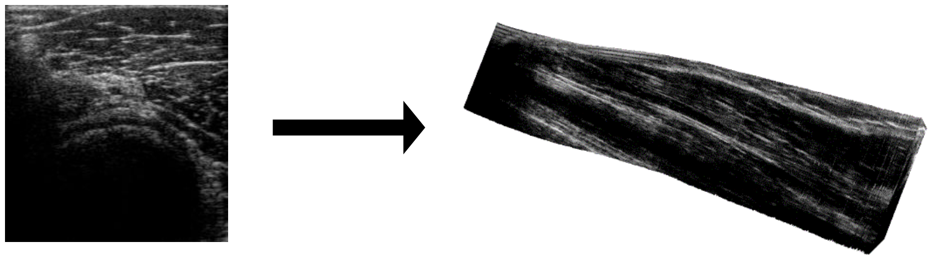

From 2D to 3D US – without External Hardware

Existing Hardware Solutions

Motorized Transducer
“wobbler”

Limited field of view
Temporal artifacts

Matrix Array
“3d probe”

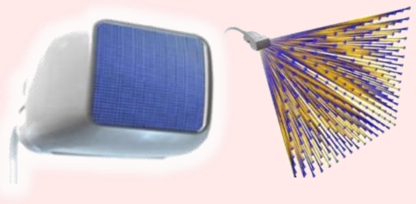

Philips xMatrix

Limited field of view
Decreased image quality

Tracking
(optical/EM)

Expensive
Not portable

Our Goal

Image-Based Reconstruction

No hardware

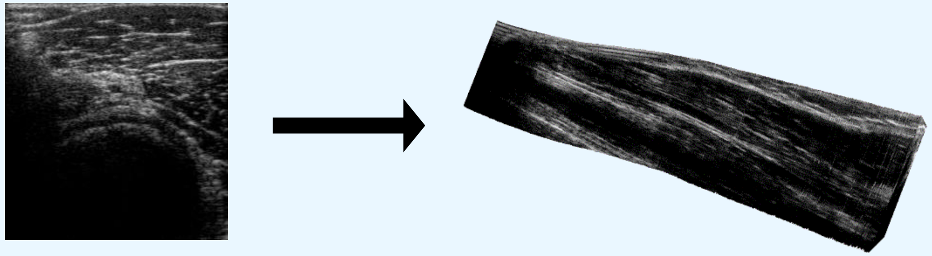

Image-based Motion Reconstruction

Frame-to-frame motion estimation

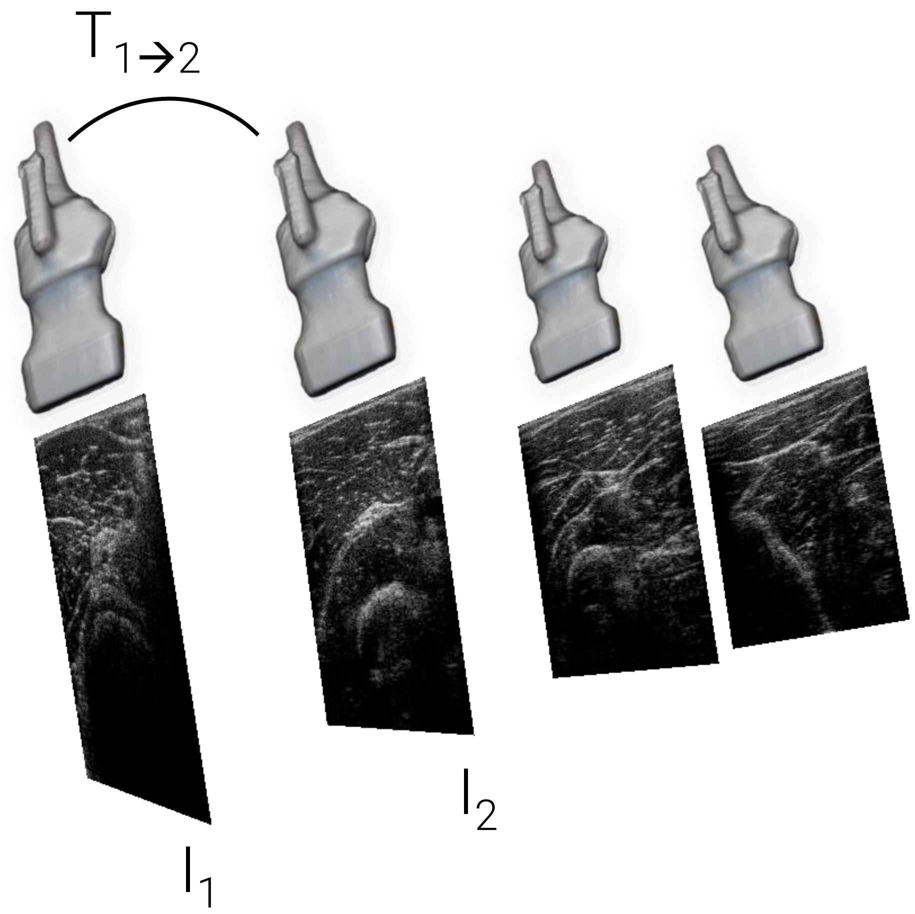

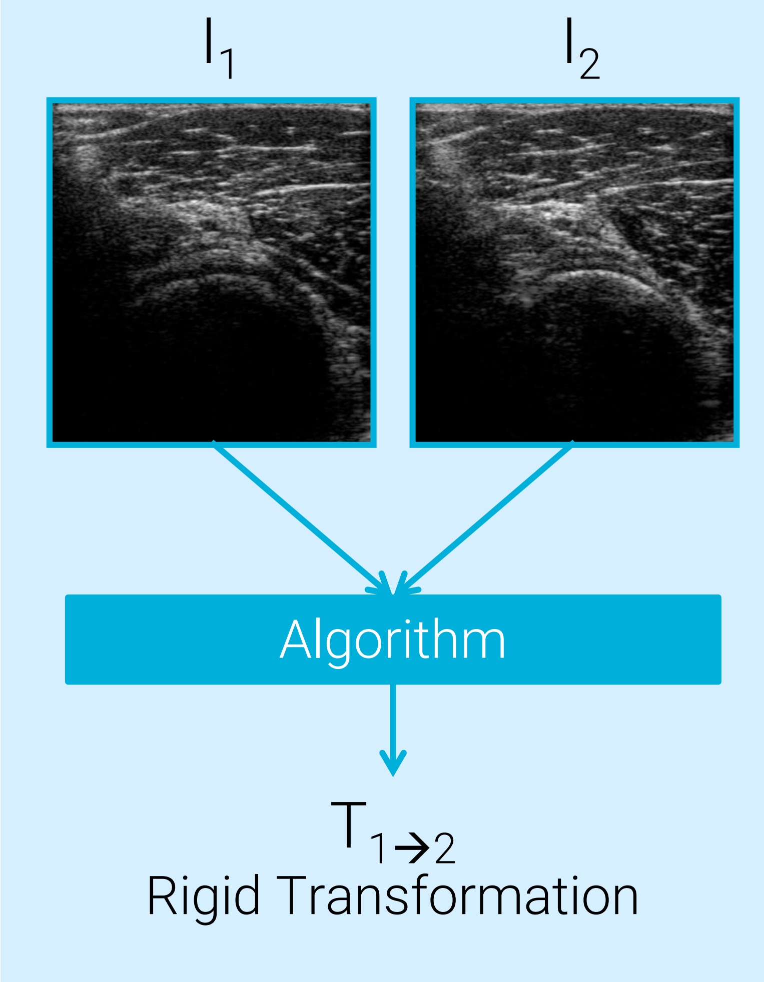

Image-based Motion Reconstruction

Frame-to-frame motion estimation

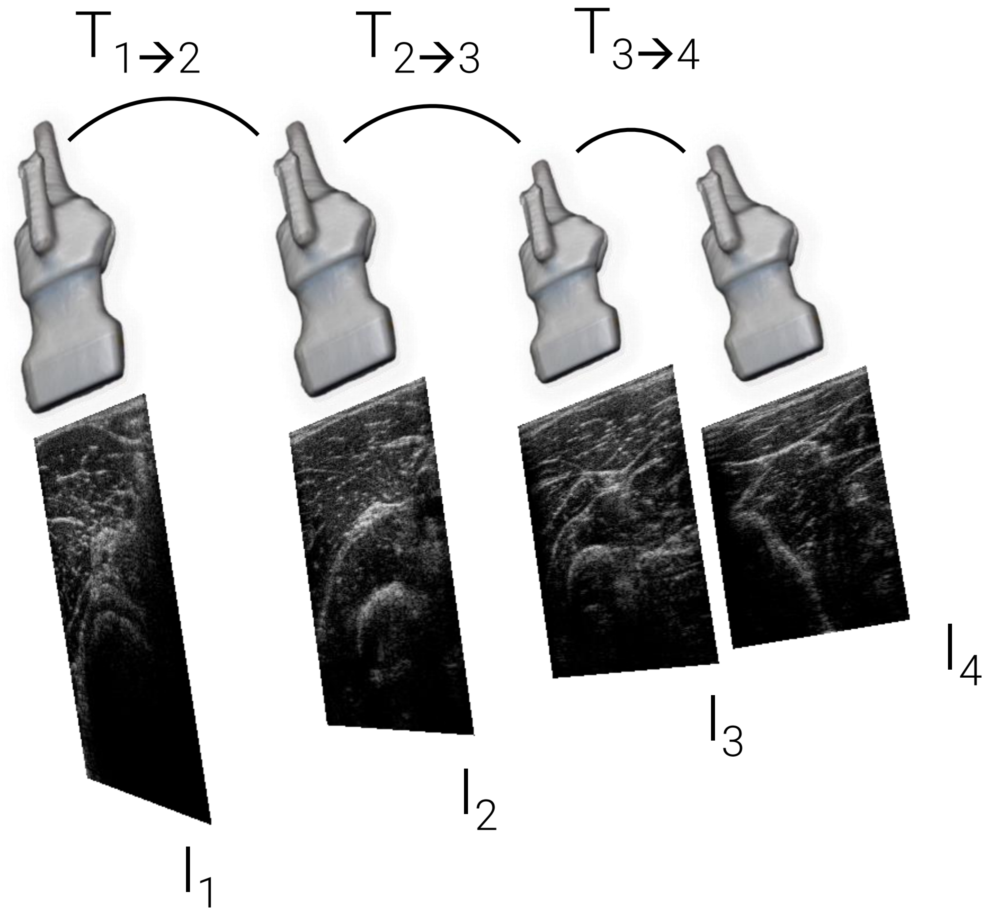

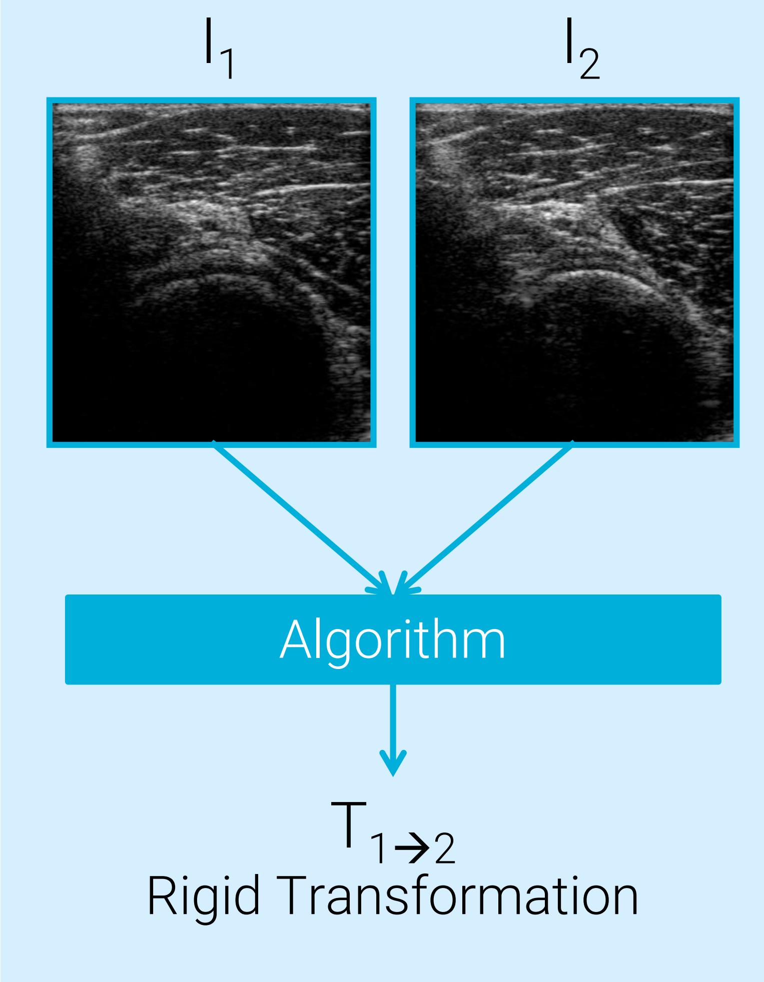

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

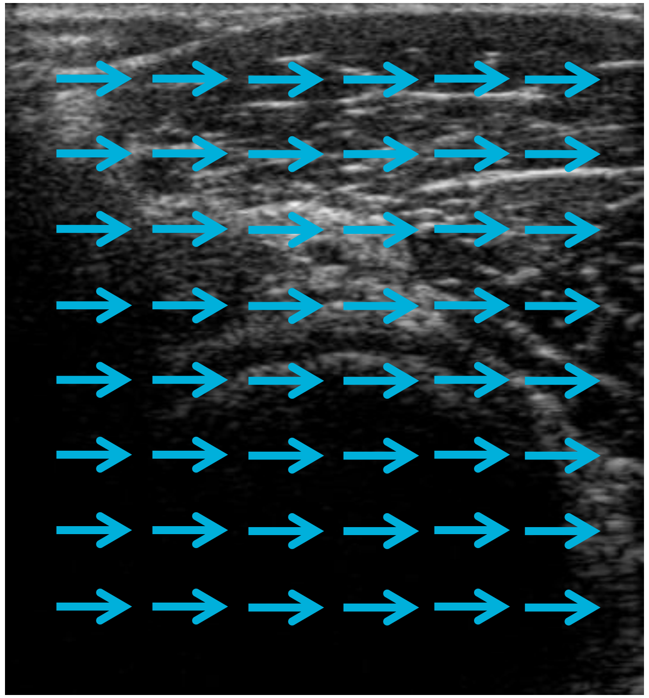

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

Out-of-plane motion
is much more difficult
because the image content changes

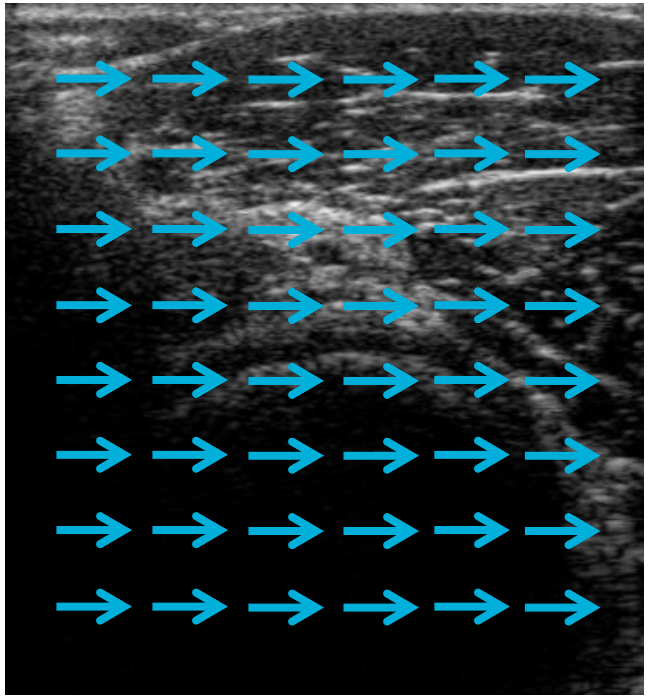

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

Out-of-plane motion
is much more difficult
because the image content changes

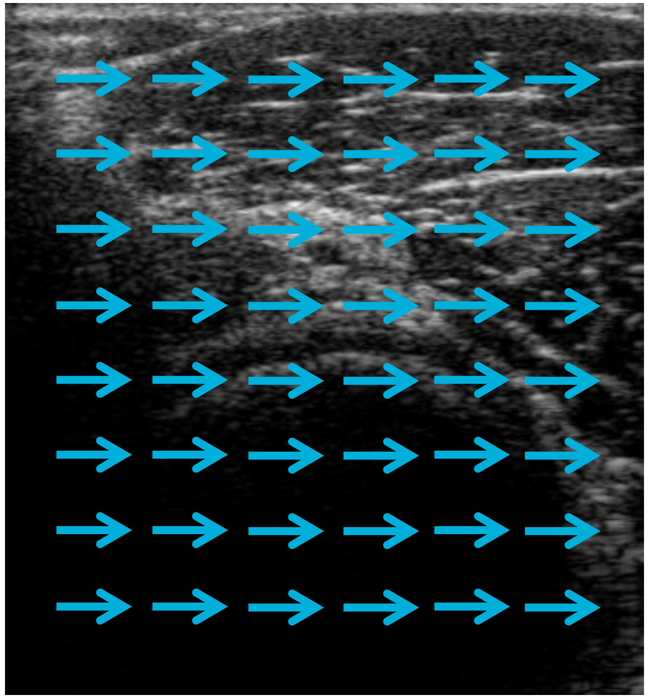

The more the content changes,
the higher the out-of-plane displacement

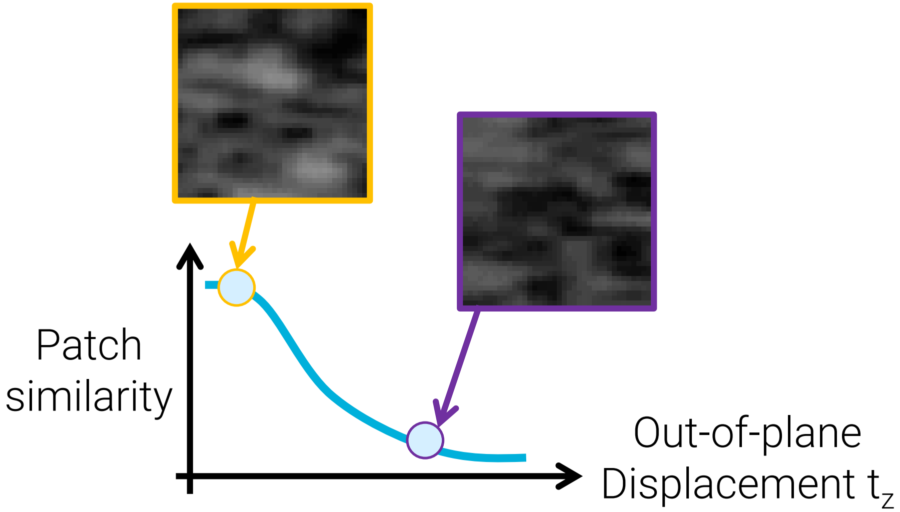

Image-based Motion Reconstruction

In-plane motion
is easy to detect
(optical flow, block matching)

Out-of-plane motion
is much more difficult
because the image content changes

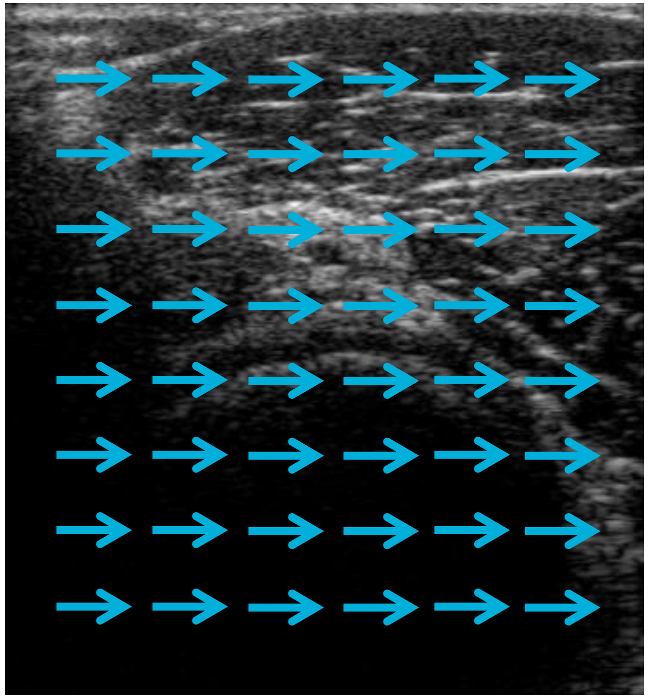

The more the content changes,
the higher the out-of-plane displacement

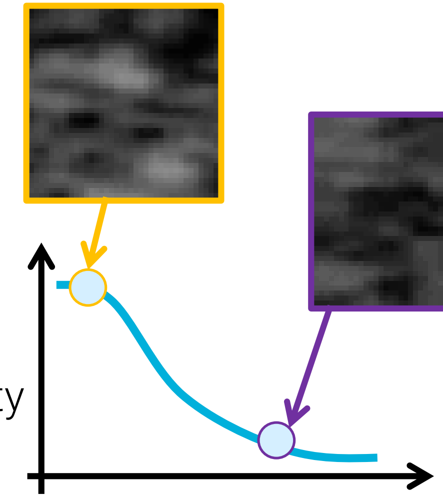

Standard approach
=
Speckle decorrelation

- Split pair of images into patches
- 2D vector field + t_z = 3D vector field
- Mask non-speckle areas
- Fit a rigid transformation

Machine Learning for Tracking Estimation

Issues of speckle decorrelation

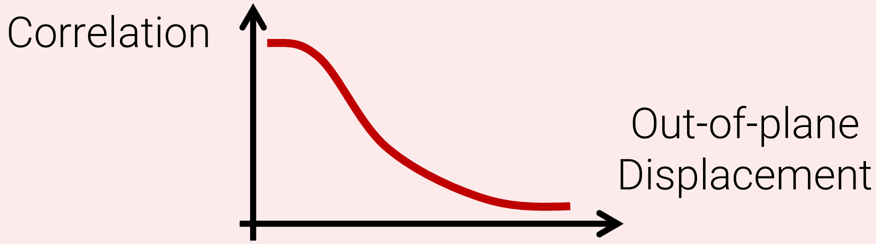

- Decorrelation is very difficult to model (depends on the tissues, on the acquisitions parameters, etc.)
- Physical model assumes Rayleigh conditions
- Errors add up through the entire pipeline (2D registration, decorrelation, transformation fitting)

Machine Learning for Tracking Estimation

Issues of speckle decorrelation

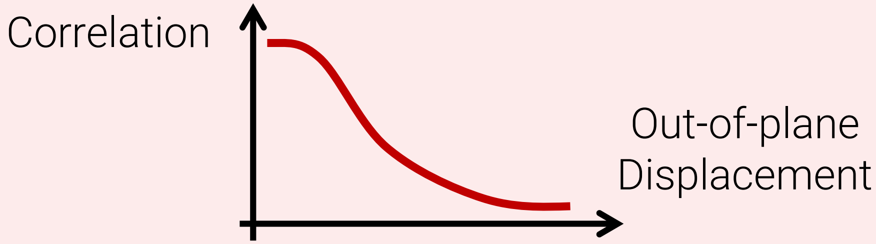

- Decorrelation is very difficult to model (depends on the tissues, on the acquisitions parameters, etc.)
- Physical model assumes Rayleigh conditions
- Errors add up through the entire pipeline (2D registration, decorrelation, transformation fitting)

Our End-to-end Approach

One model to solve the whole problem

pair of images → transformation parameters

Convolutional Neural Network

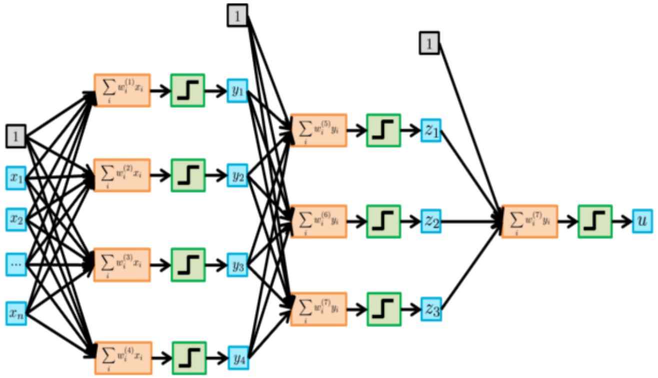

Training Data Acquisition

No need for manual labeling
We just need to acquire a lot of tracked sweeps
(but calibration must be super accurate)

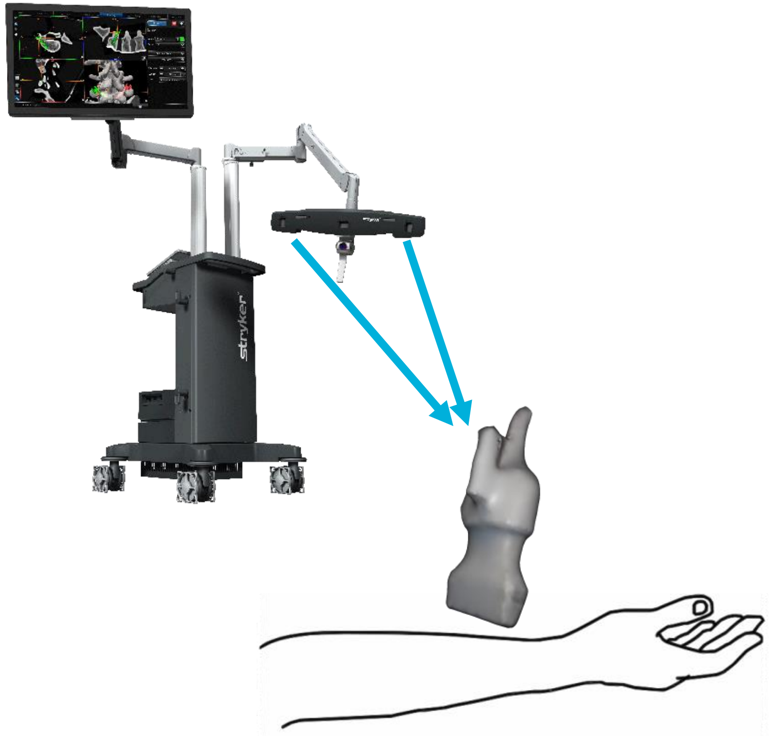

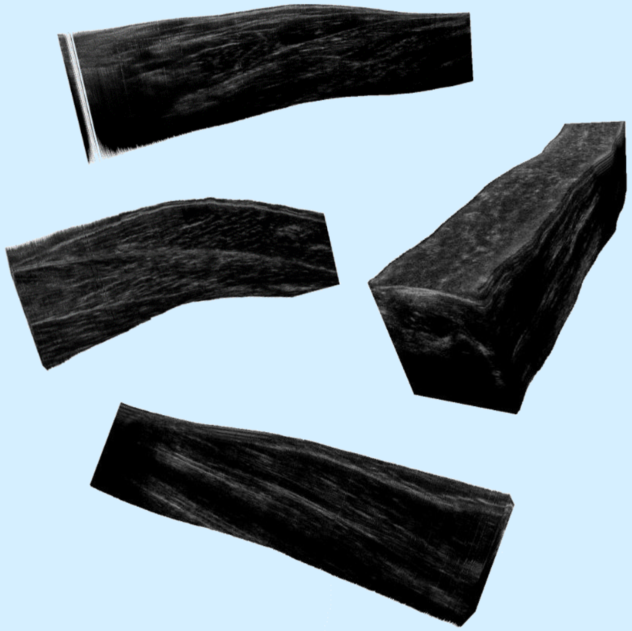

800 sweeps (400k frames)

Neural Network Architecture

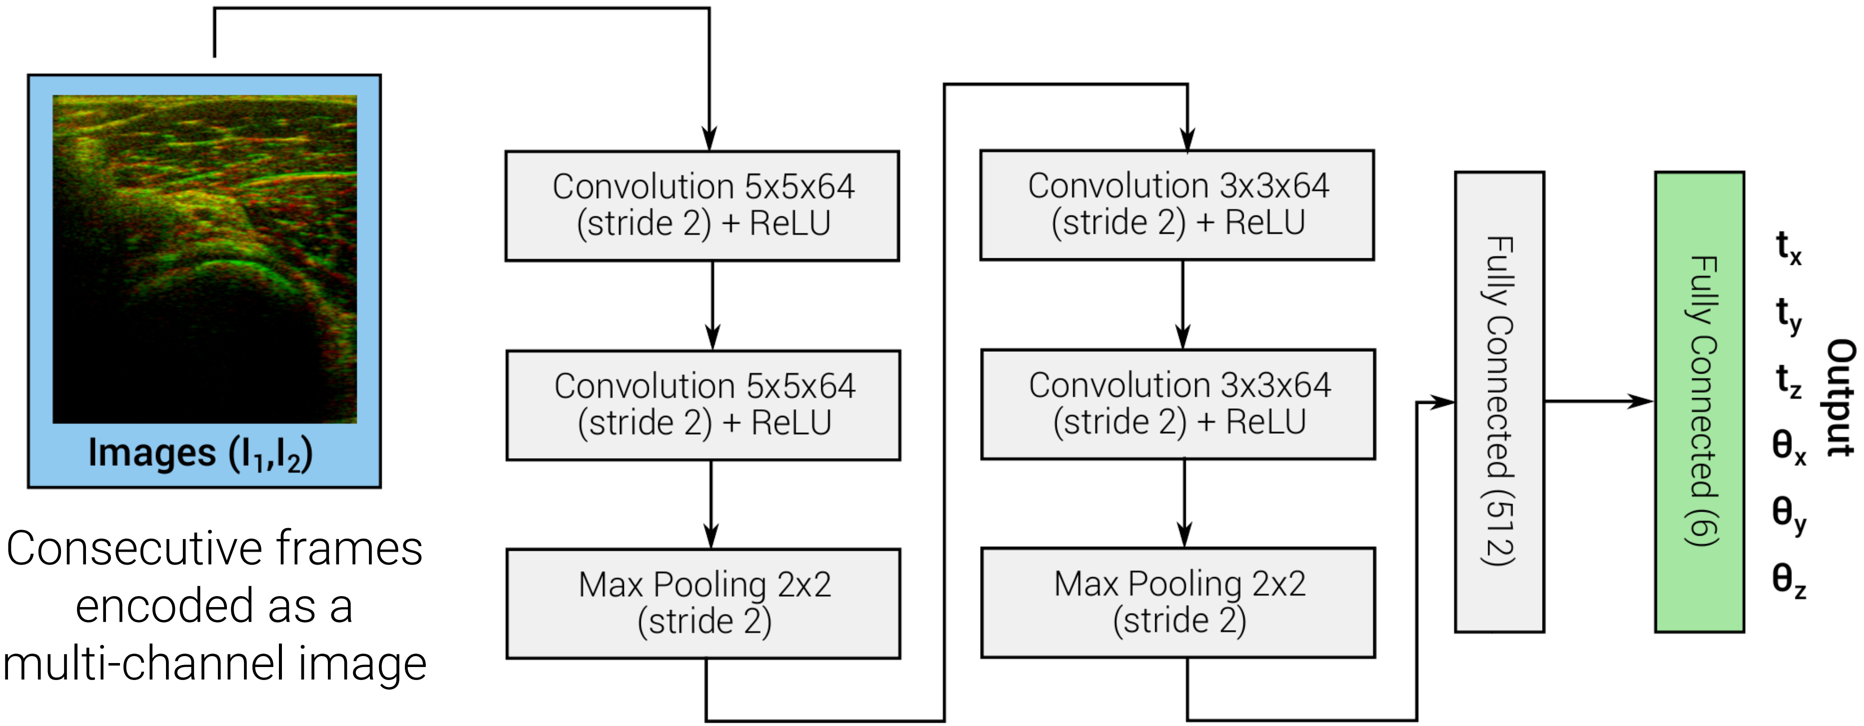

Consecutive frames encoded as a multi-channel image

Regression
L2 loss

3 translation + 3 rotation parameters
(probe motion between the 2 images)

Trick #1: Use the optical flow

Pre-compute the optical flow (in-plane motion)
and use it as additional channel

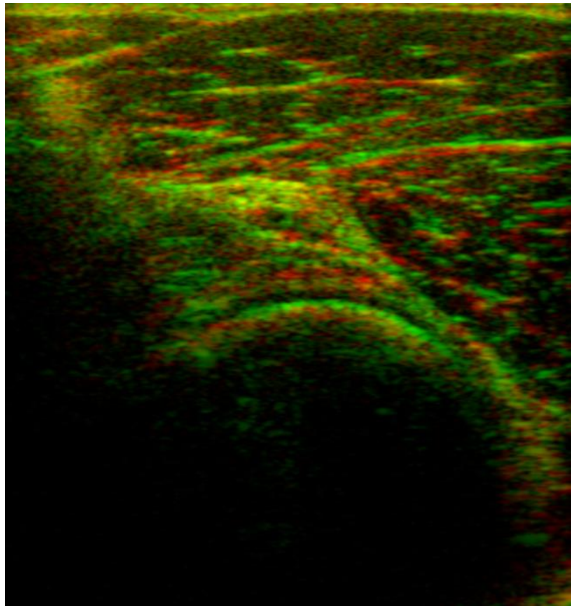

2-channel input
2 ultrasound images

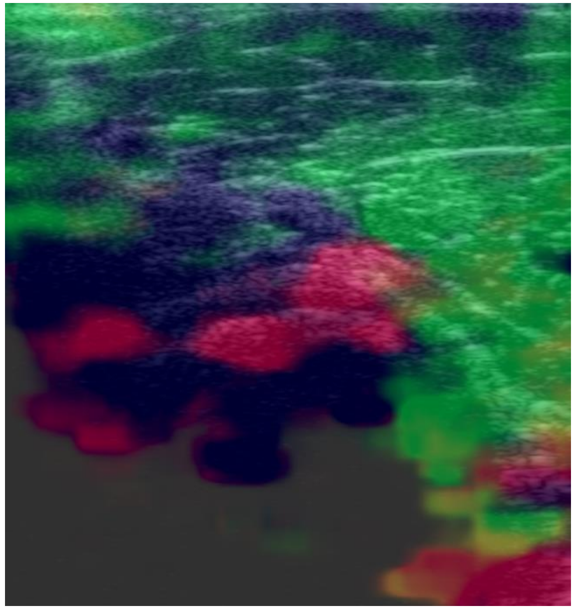

4-channel input
2 ultrasound images + 2D vector field

Trick #1: Use the optical flow

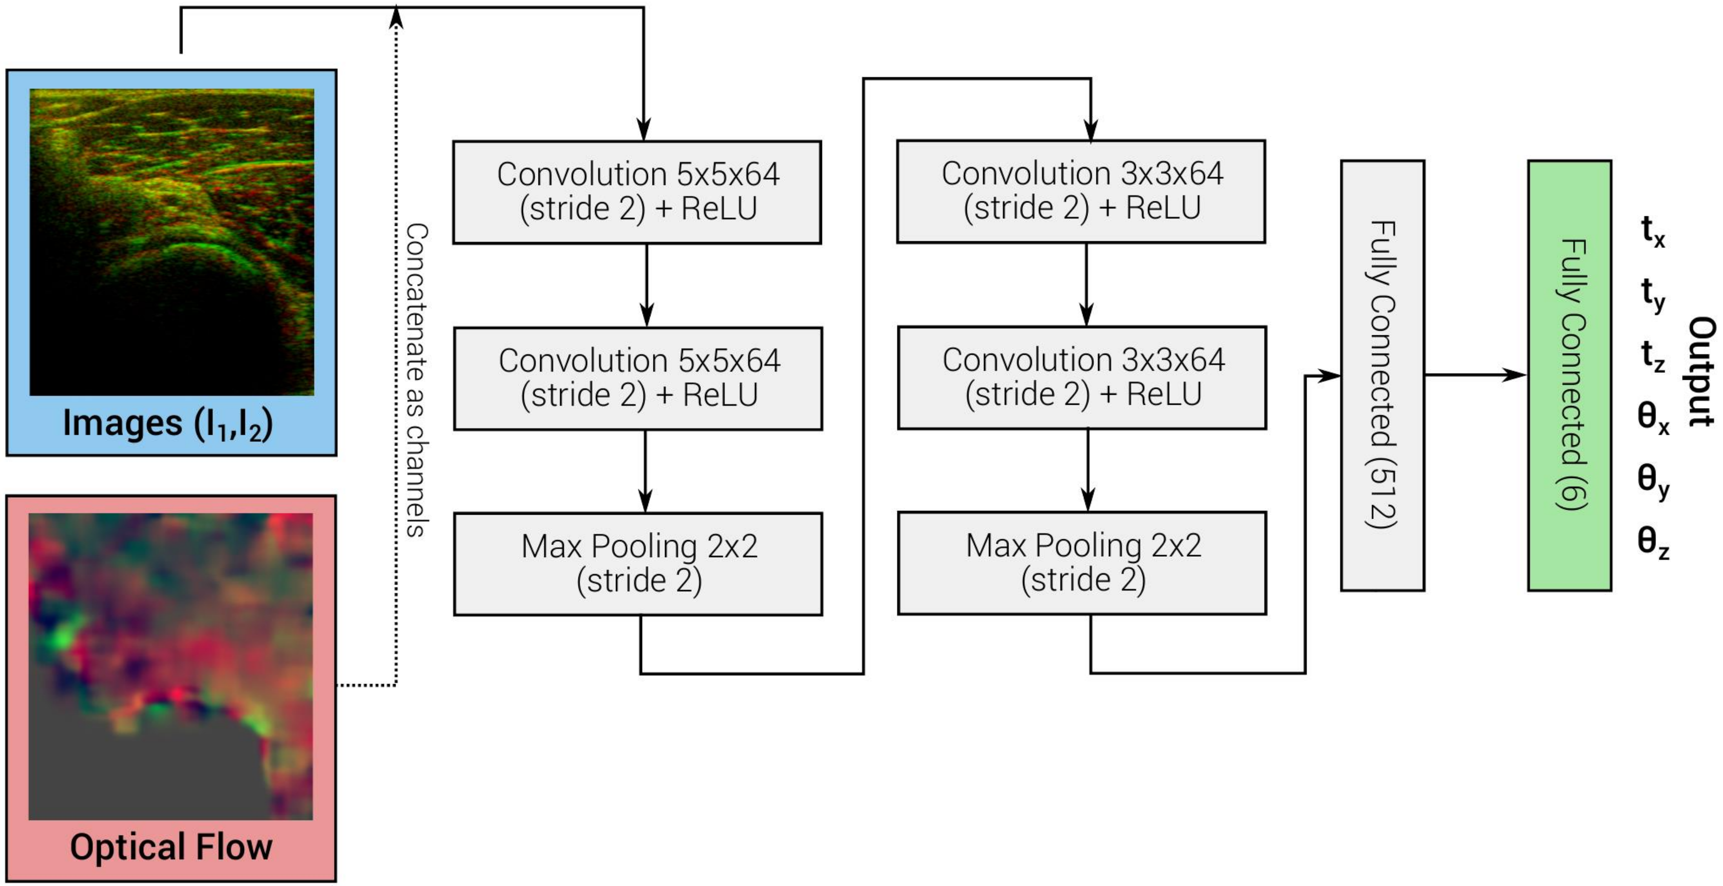

Trick #2: Use the Inertial Measurement Unit

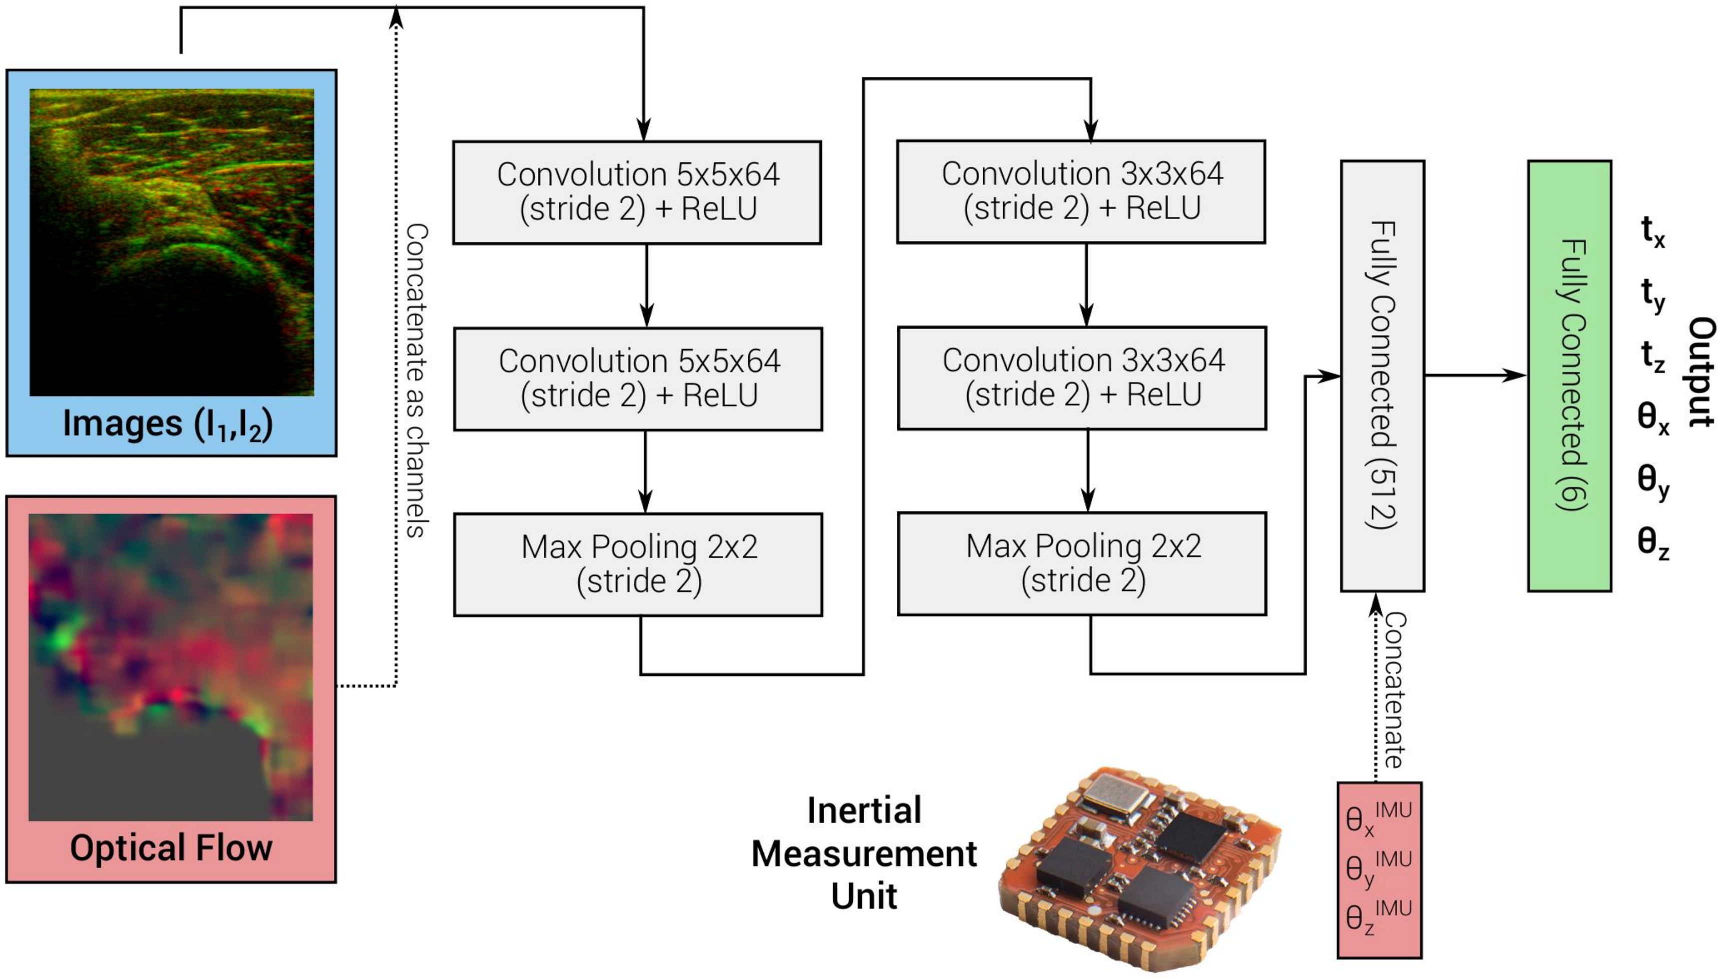

3D Reconstructions with IMU

Reconstruction of sweeps with strong motions and rotations

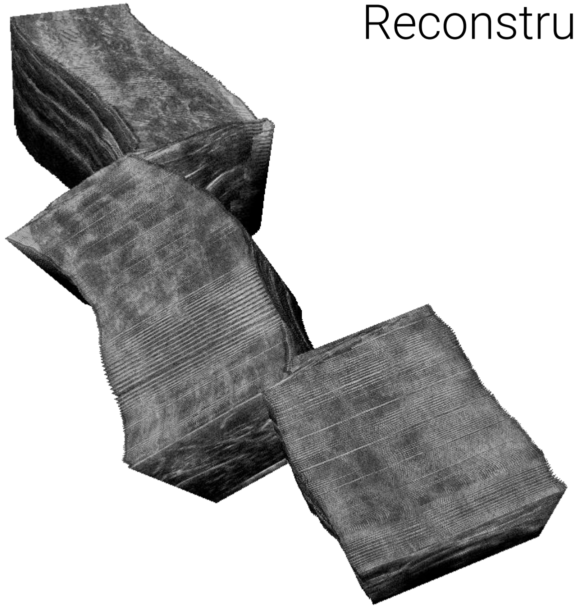

3D Reconstructions with IMU

Reconstruction of sweeps with strong motions and rotations

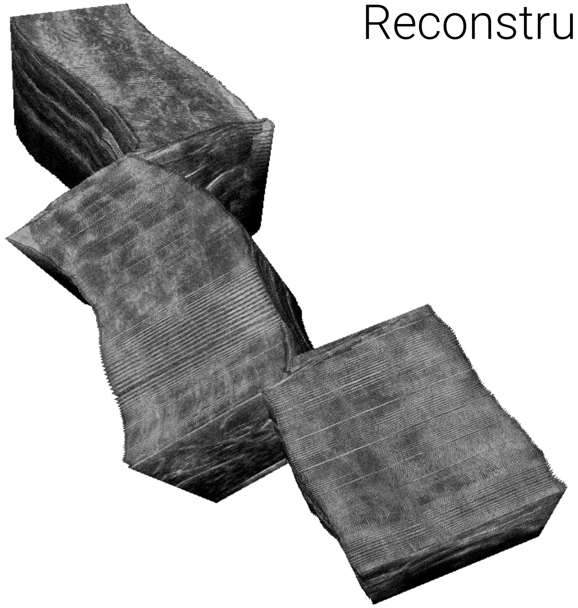

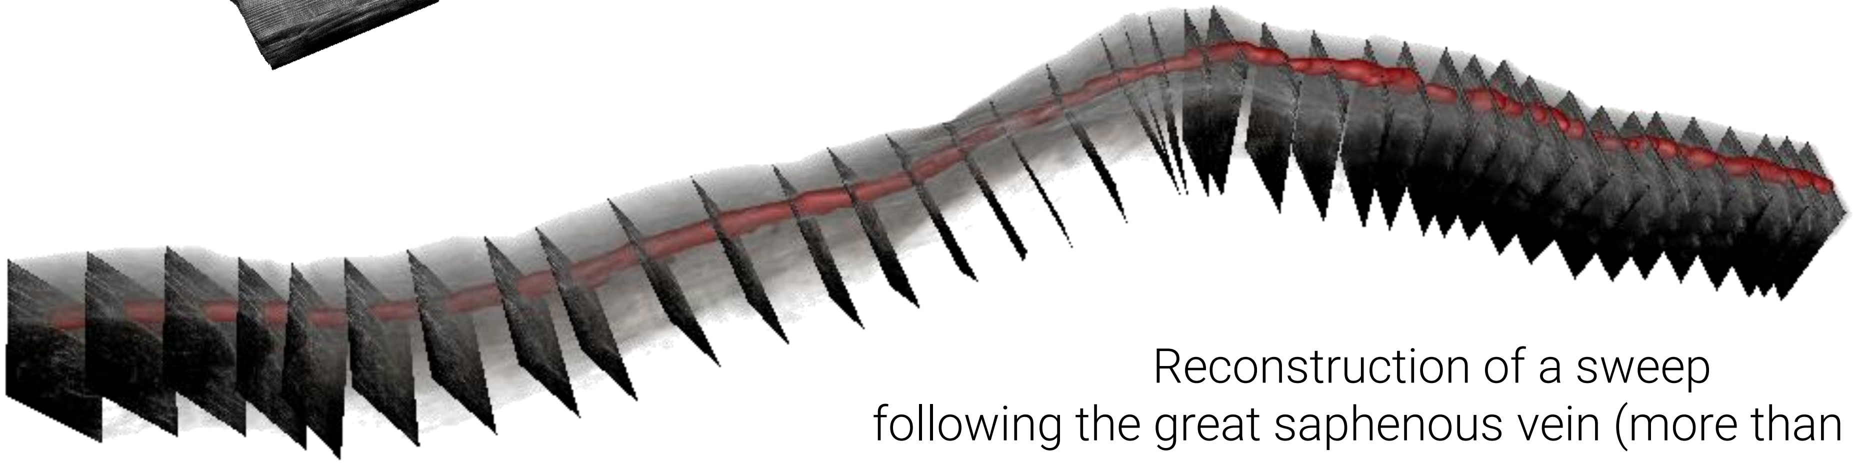

Reconstruction of a sweep
following the great saphenous vein (more than 60cm)

Quantitative Results

Accuracy study conducted on 800 US sweeps on various anatomies

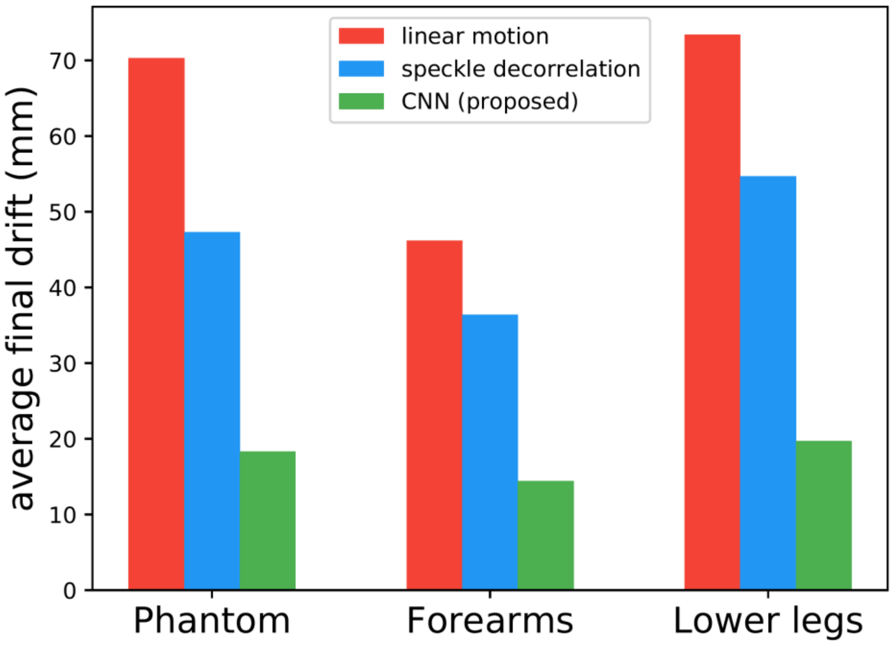

Outperforms state-of-the-art methods

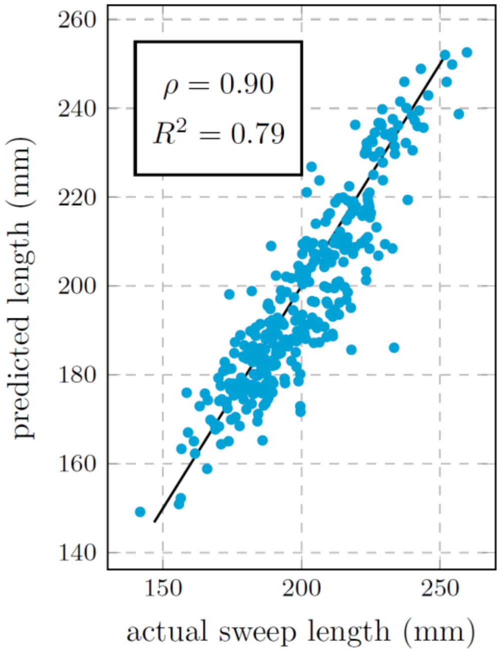

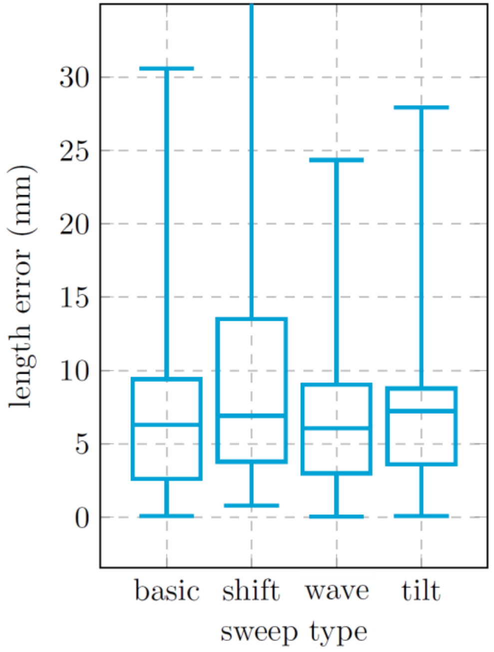

Median drift of 5% over long sweeps

More Quantitative Results

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

ELSEVIER

3D freehand ultrasound without external tracking using deep learning

Raphael Prevost[a,*], Mehrdad Salehi[a,b], Simon Jagoda[a], Navneet Kumar[a], Julian Sprung[c], Alexander Ladikos[a], Robert Bauer[c], Oliver Zettinig[a], Wolfgang Wein[a]

[a] ImFusion GmbH, Agnes-Pockels-Bogen 1, Munich, Germany
[b] Computer Aided Medical Procedures (CAMP), TU Munich, Munich, Germany
[c] Piur Imaging GmbH, Vienna, Austria

A R T I C L E I N F O

Article history:
Received 8 February 2018
Revised 5 June 2018
Accepted 6 June 2018
Available online 15 June 2018

Keywords:
3D freehand ultrasound
Deep learning
Motion estimation
Inertial measurement unit

A B S T R A C T

This work aims at creating 3D freehand ultrasound reconstructions from 2D probes with image-based tracking, therefore not requiring expensive or cumbersome external tracking hardware. Existing model-based approaches such as speckle decorrelation only partially capture the underlying complexity of ultrasound image formation, thus producing reconstruction accuracies incompatible with current clinical requirements. Here, we introduce an alternative approach that relies on a statistical analysis rather than physical models, and use a convolutional neural network (CNN) to directly estimate the motion of successive ultrasound frames in an end-to-end fashion. We demonstrate how this technique is related to prior approaches, and derive how to further improve its predictive capabilities by incorporating additional information such as data from inertial measurement units (IMU). This novel method is thoroughly evaluated and analyzed on a dataset of 800 in vivo ultrasound sweeps, yielding unprecedentedly accurate reconstructions with a median normalized drift of 5.2%. Even on long sweeps exceeding 20 cm with complex trajectories, this allows to obtain length measurements with median errors of 3.4%, hence paving the way toward translation into clinical routine.

All papers and references available on www.imfusion.com

Live Demo
Carotid Reconstruction

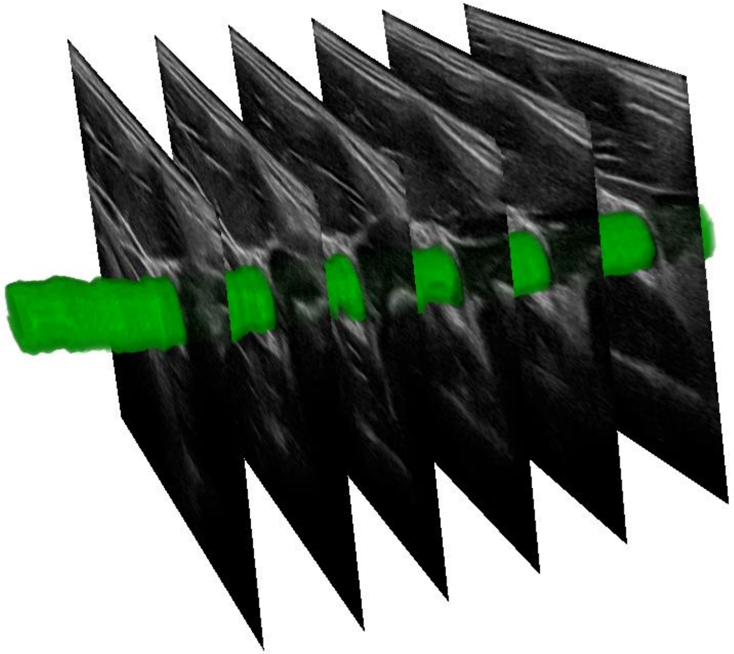

in partnership with

Cephasonics Ultrasound Solutions

What if one sweep is not enough?

What if one sweep is not enough?

What if one sweep is not enough?

What if one sweep is not enough?

Anatomical structures do not match
because of compression

What if one sweep is not enough?

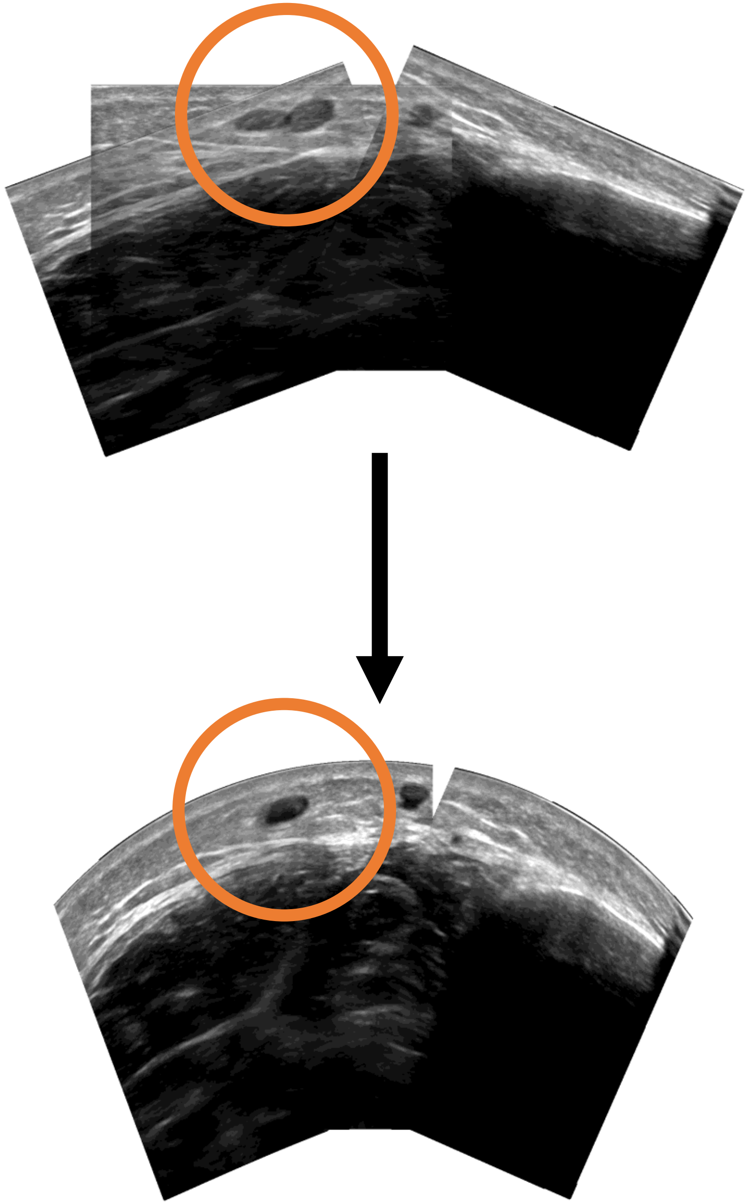

Anatomical structures do not match
because of compression

Decompression Model

- Skin surface locally modeled as a circle

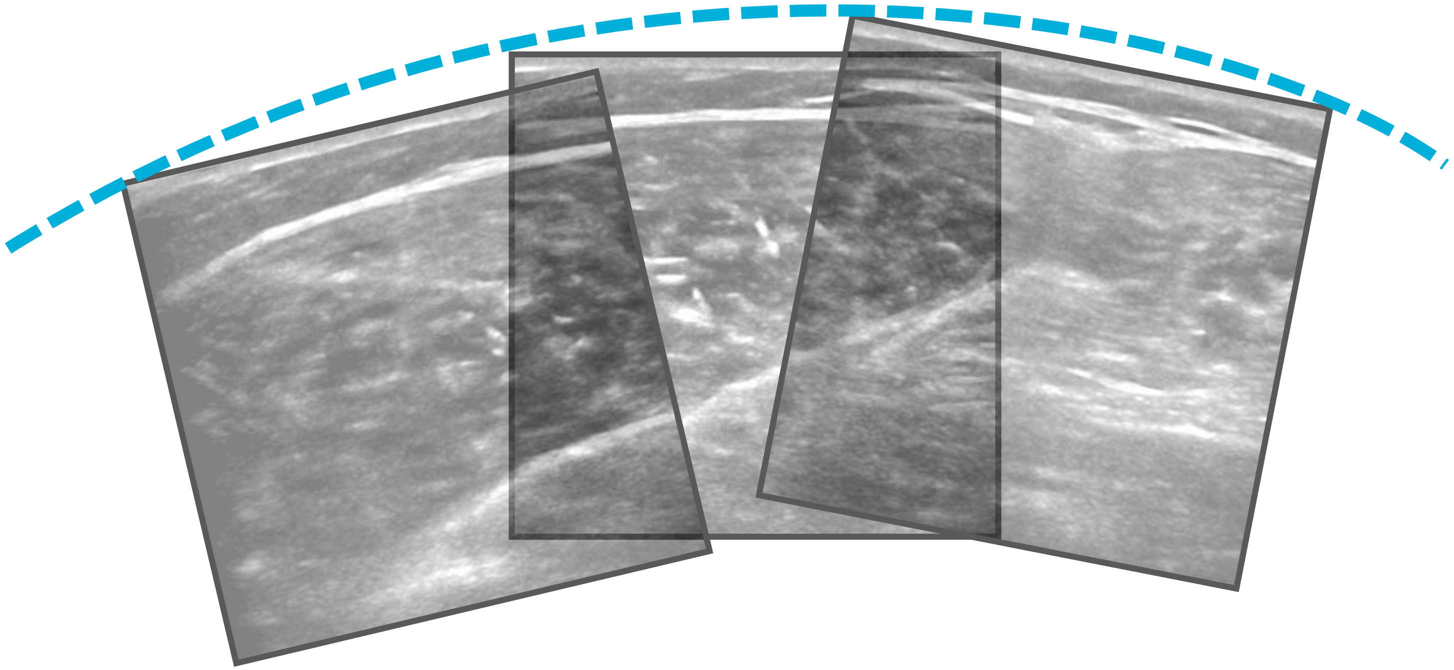

Decompression Model

- Skin surface locally modeled as a circle

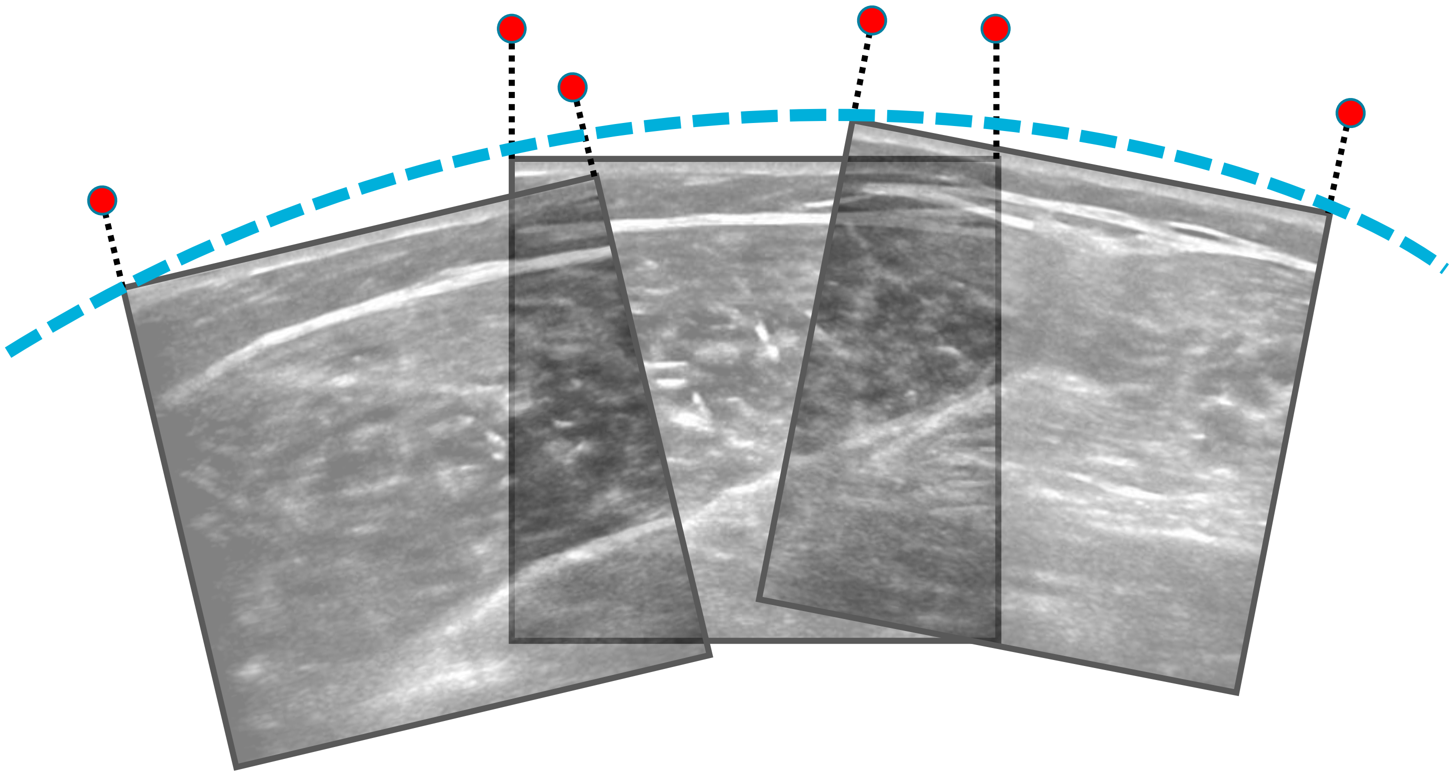

Decompression Model

- Skin surface locally modeled as a circle

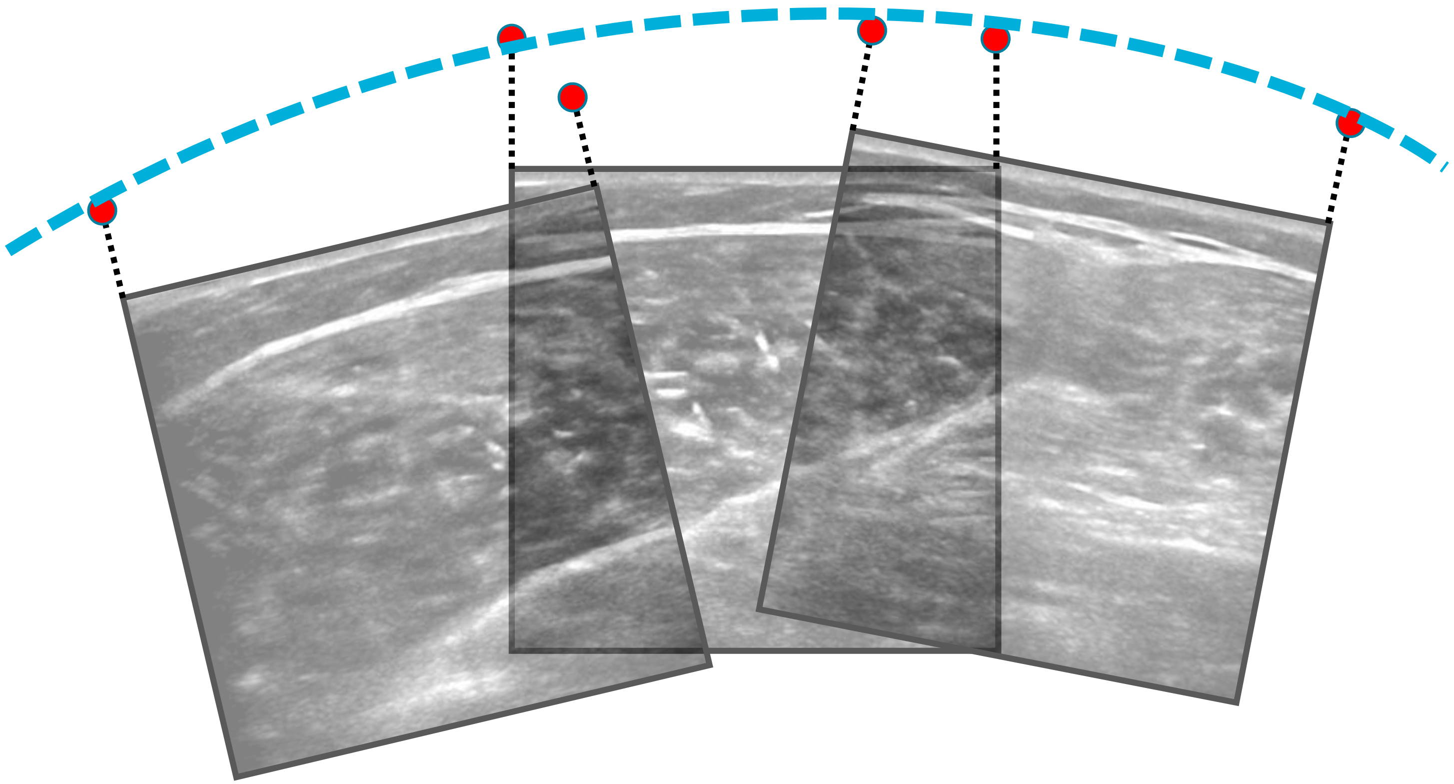

Decompression Model

- Skin surface locally modeled as a circle
- Displacements are optimized by maximizing image similarity in the overlapping regions

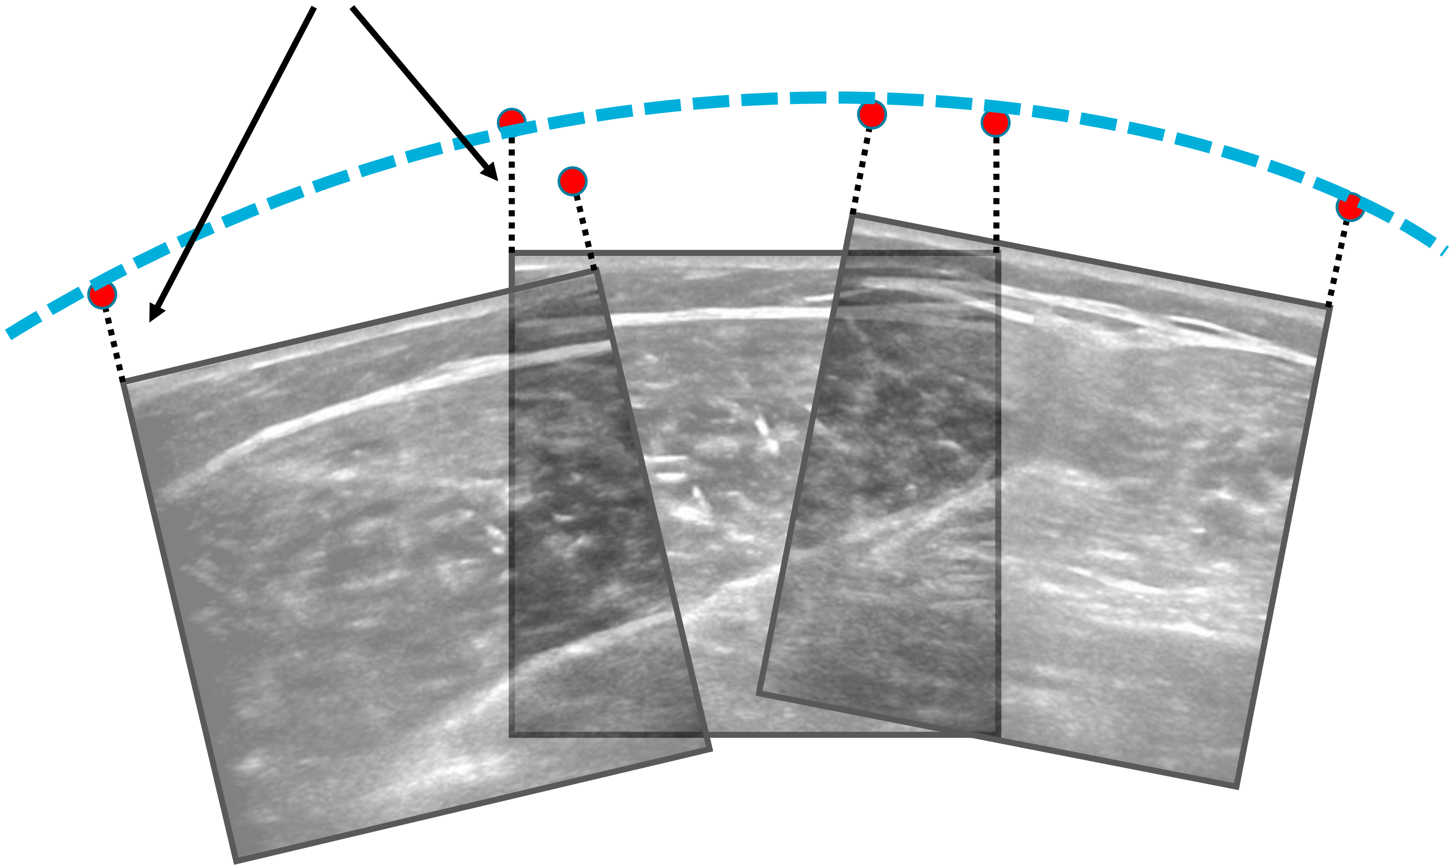

Decompression Model

- Skin surface locally modeled as a circle
- Displacements are optimized by maximizing image similarity in the overlapping regions

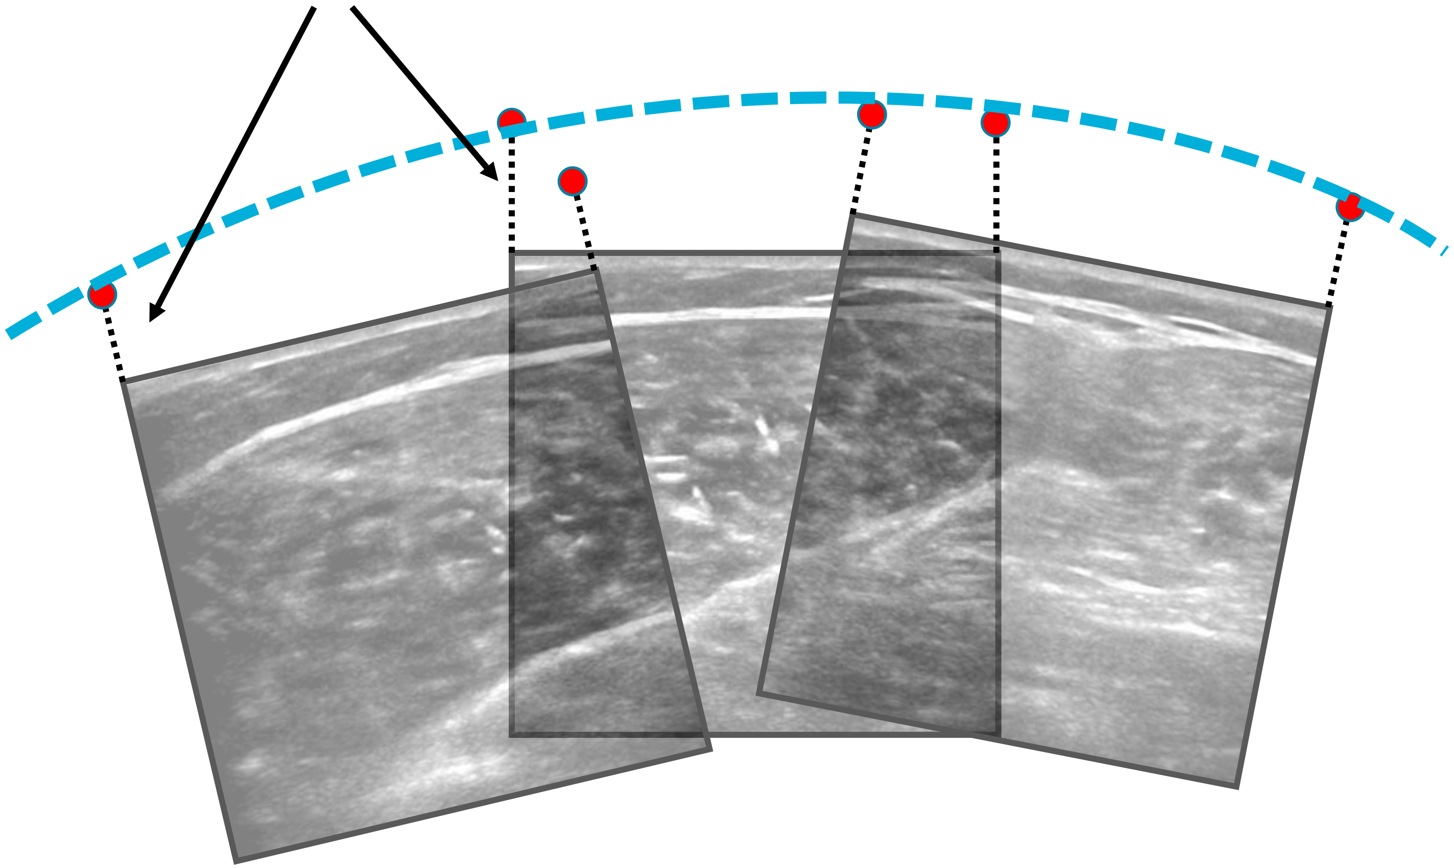

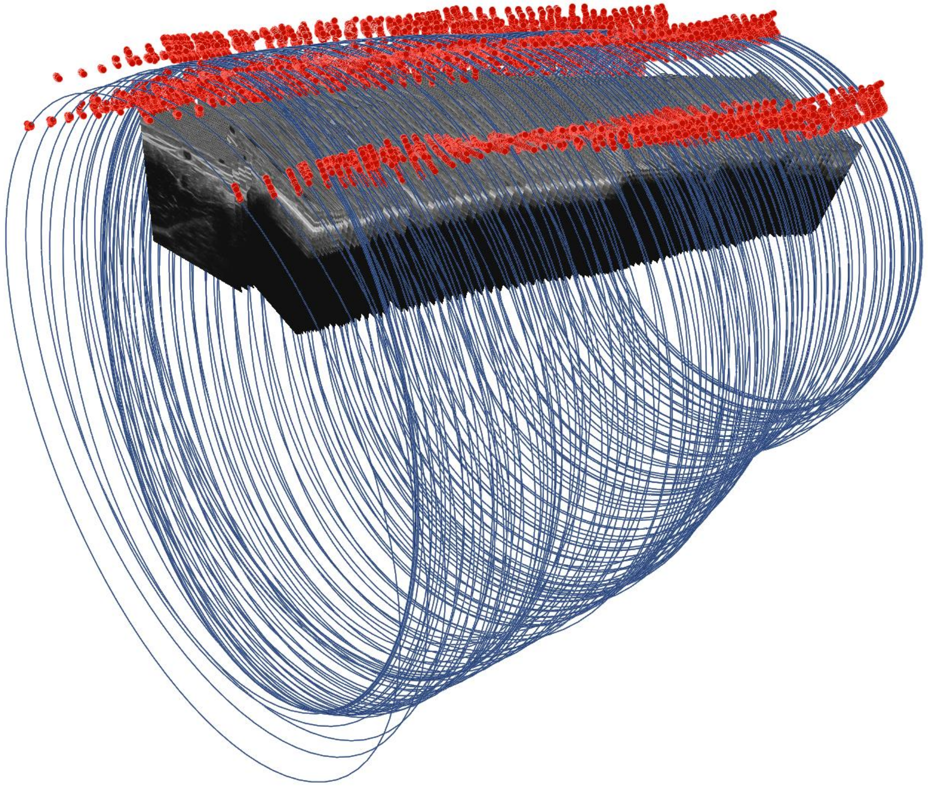

Multi-scan Decompression Algorithm

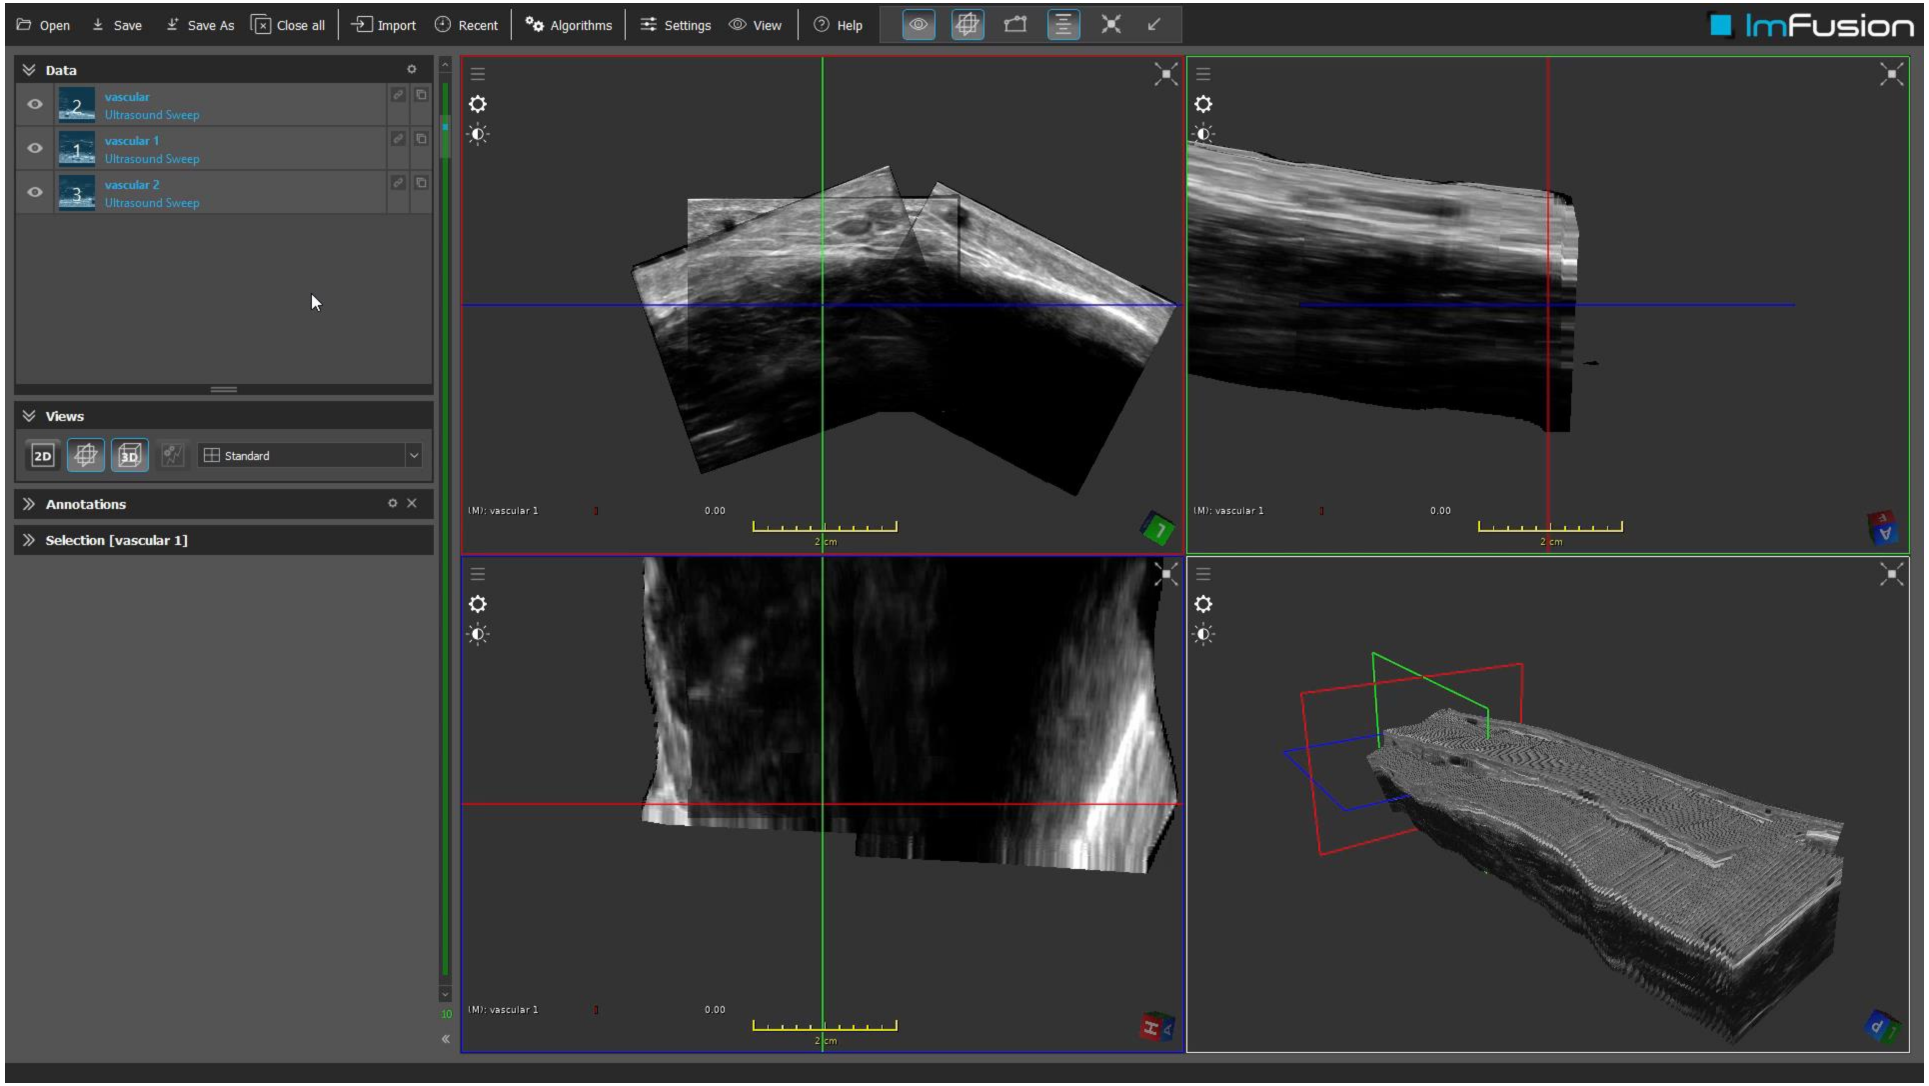

Wide Field-of-View Reconstruction

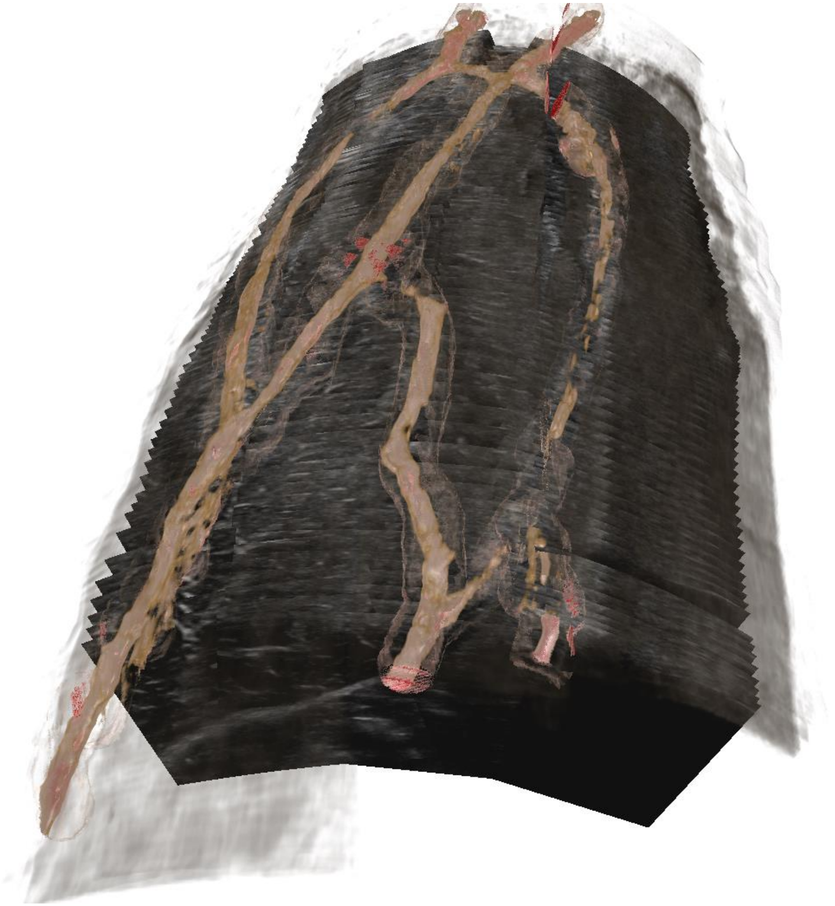

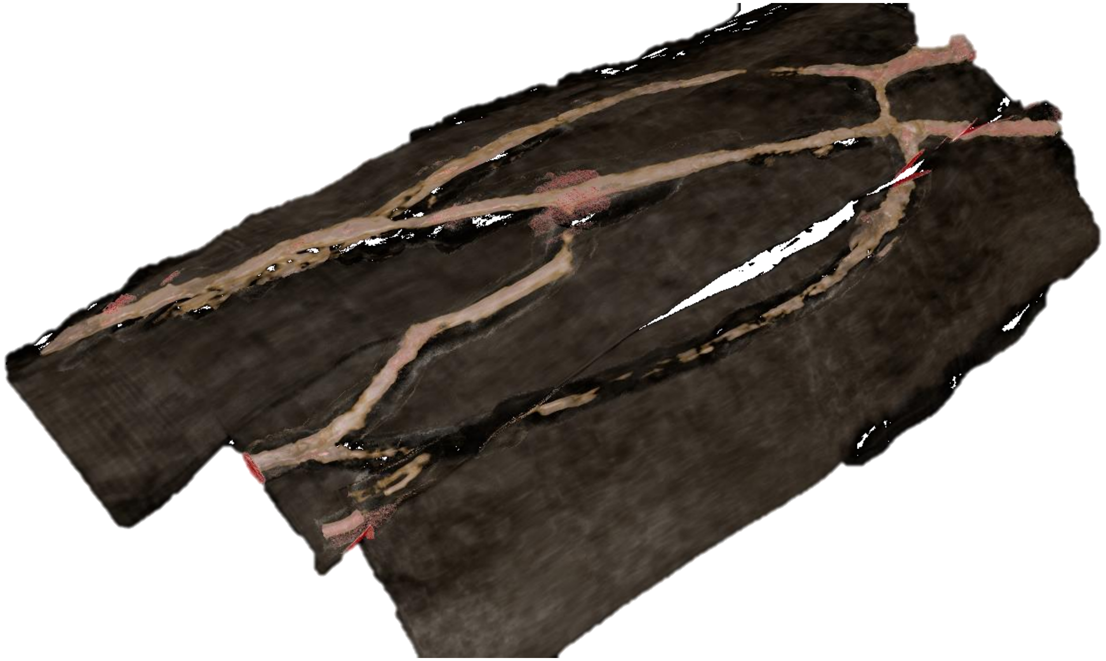


EG VCBM 2018

8th EUROGRAPHICS Workshop on
VISUAL COMPUTING FOR BIOLOGY AND MEDICINE – Granada (Spain), 20-21 September 2018

Best Paper Award

presented to

Christian Schulte zu Berge, Wolfgang Wein, Mehrdad Salehi, Frederik Bender

for their work

Ultrasound Decompression for Large Field-of-View Reconstructions

EG VCBM 2018

VCBM 2018 General Chairs: F. Javier Melero, Anna Vilanova

Chairs Best Paper Committee: Noeska Smit, Barbora Kozlikova, Steffen Oeltze-Jafra

Schulte zu Berge et al.

Ultrasound Decompression for Large Field-of-View Reconstructions
VCBM 2018

Conclusion

The Future of Ultrasound Imaging

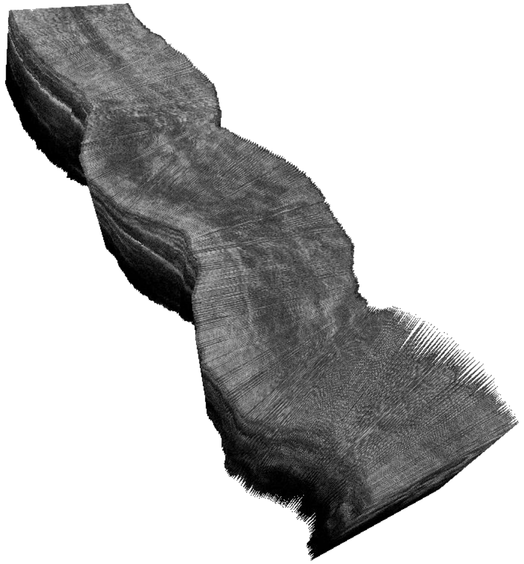

Let’s recap

Let’s recap

- Ultrasound acquisition can be made easier and less tedious

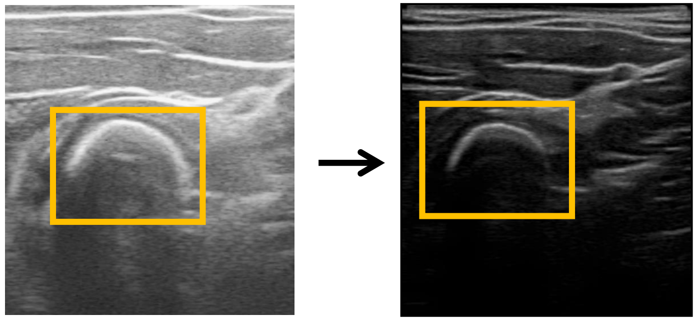

Auto-tuning of the parameters

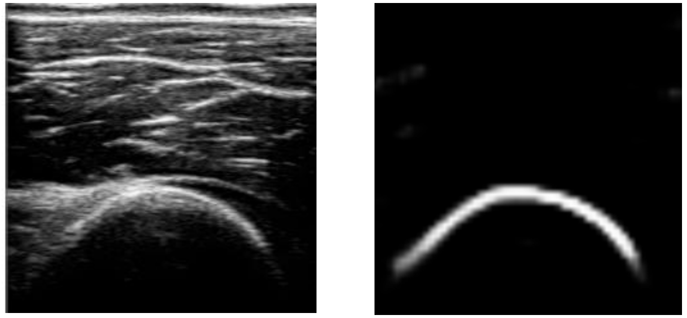

Real-time anatomy recognition

Trackingless 3D Reconstruction

Let's recap

- Ultrasound acquisition can be made easier and less tedious

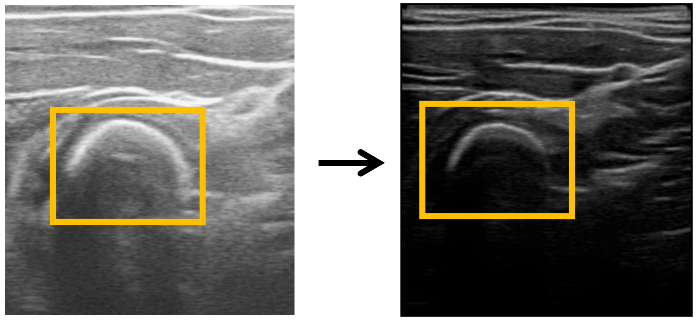

Auto-tuning of the parameters

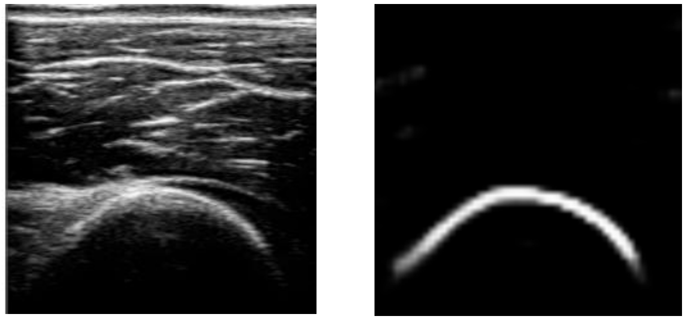

Real-time anatomy recognition

Trackingless 3D Reconstruction

- Ultrasound improves both surgery workflows and diagnostics/monitoring

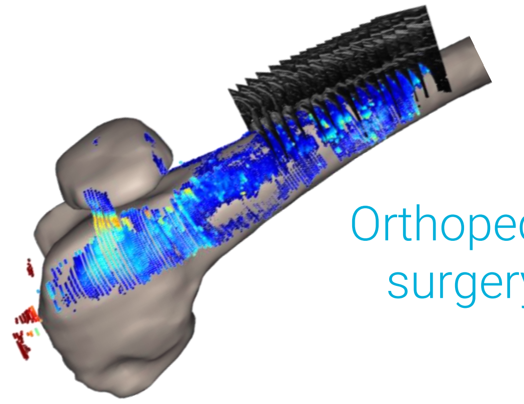

Orthopedic surgery

Neuro surgery

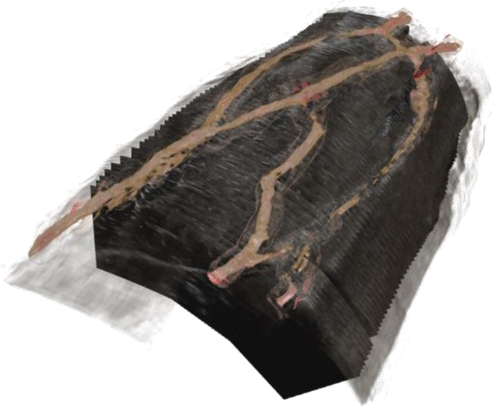

Vascular imaging

Let's recap

- Ultrasound acquisition can be made easier and less tedious

With AI + GPU computing + advanced algorithms,
US becomes more accessible and create new applications
… maybe even replace other modalities in the long run

Orthopedic surgery

Neuro surgery

Vascular imaging

ImFusion Suite: The ideal platform for R&D

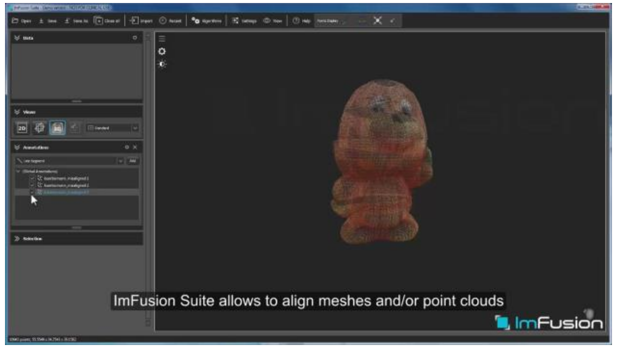

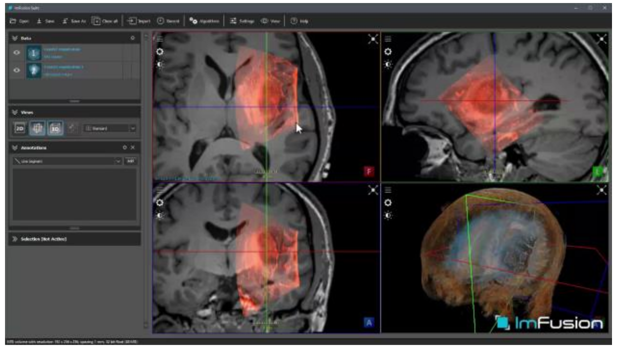

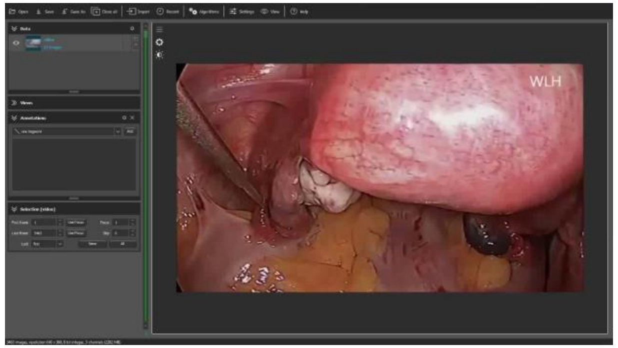

ImFusion Suite: The ideal platform for R&D

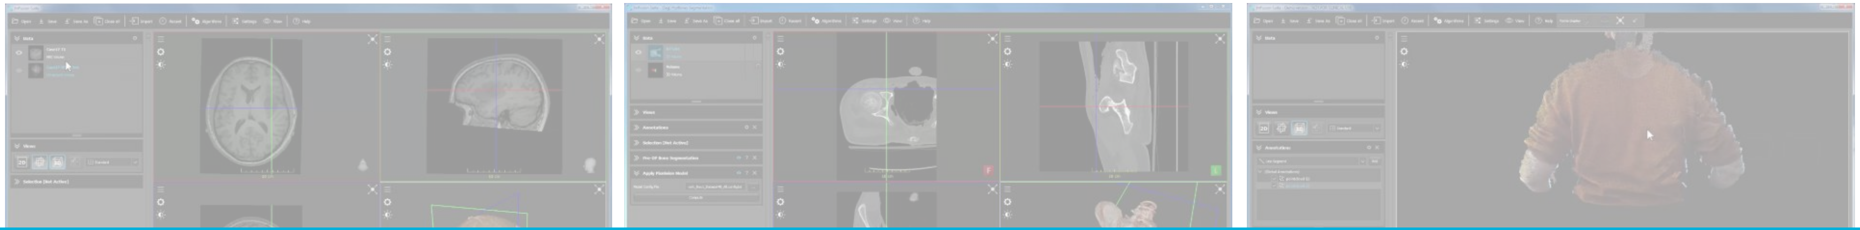

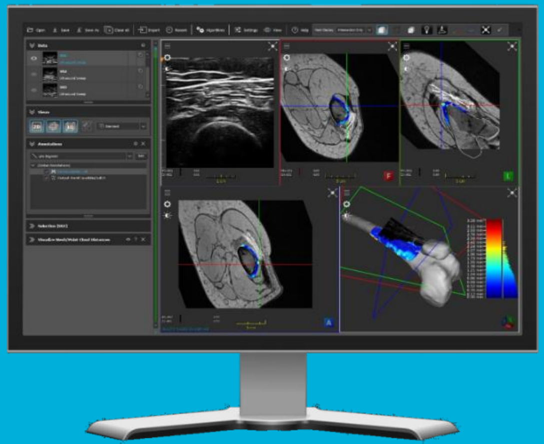

Download the ImFusion Suite demo
www.imfusion.com

Image Visualization,
Segmentation, Registration,
Mesh/Point Cloud Processing,
… and more!

ImFusion x NVIDIA

NVIDIA Clara initiative for transparent access to accelerated computing (closer to the sensor/raw data for certain applications & high-end systems vs. in the cloud for point-of-care ultrasound)

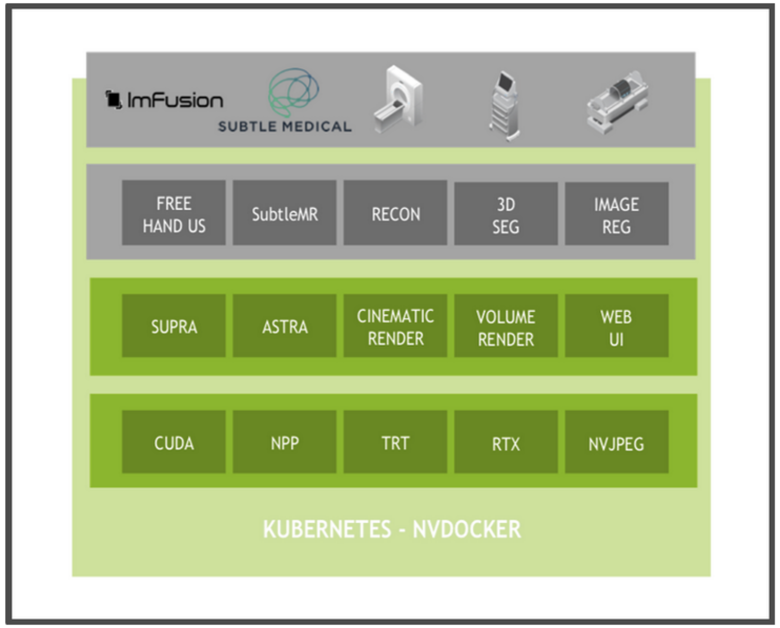

NVIDIA Clara SDK

Clara Developers

source: https://developer.nvidia.com/clara

ImFusion SDK x CLARA Rendering Server

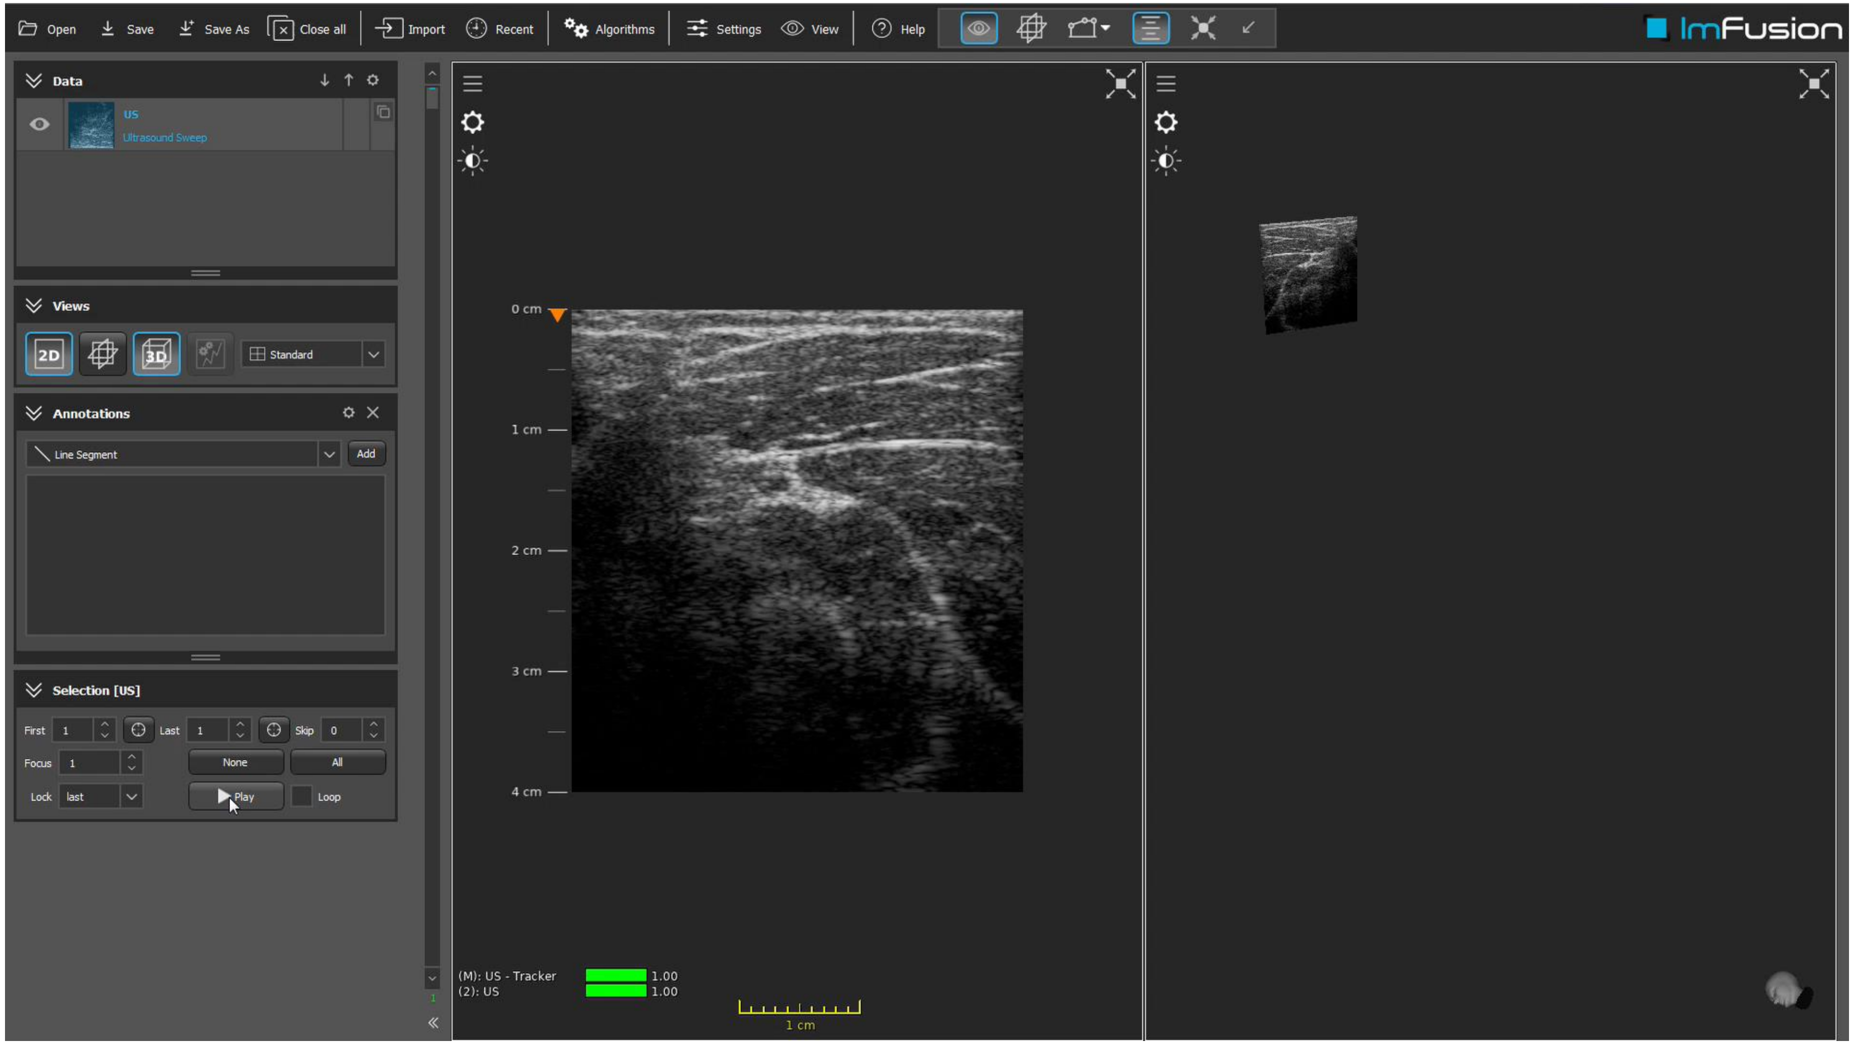

Thank you!

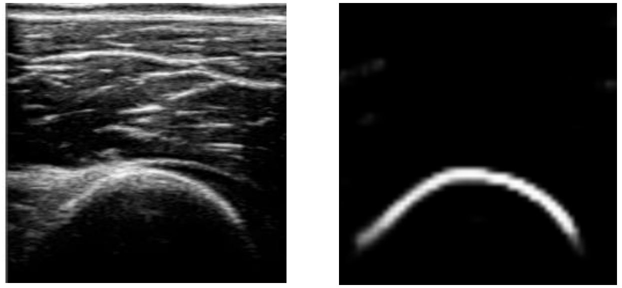

Real-time anatomy detection

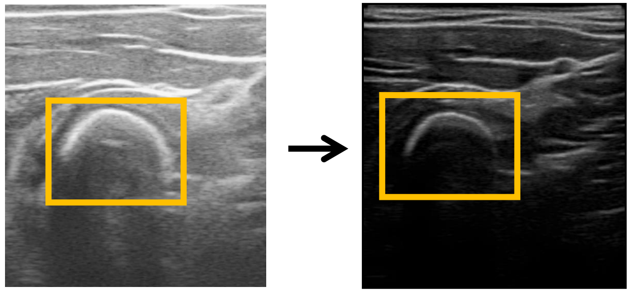

Auto-tuning of the parameters

Tracking-less 3D ultrasound

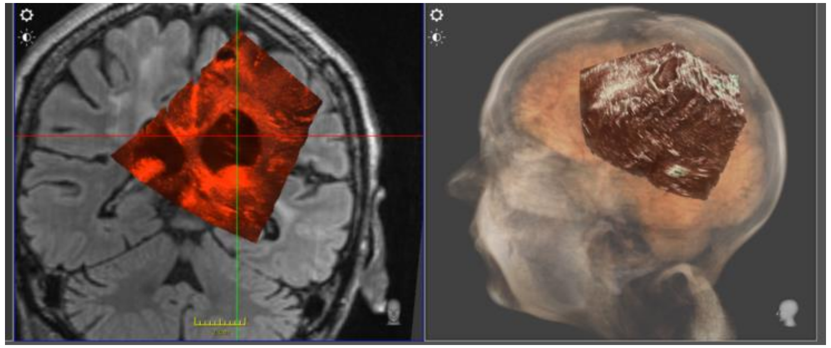

Multi-modal registration

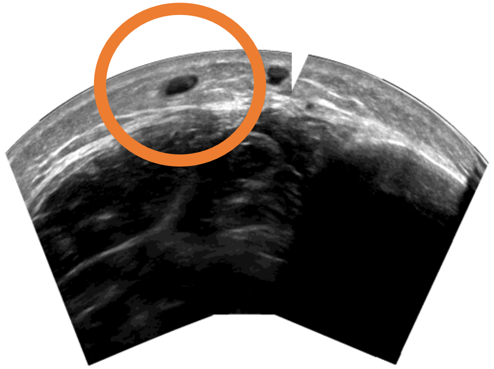

Decompression & Stitching

Raphael Prevost
E-Mail: prevost@imfusion.com
Web: www.imfusion.com