Developing and deploying scalable, efficient, and accurate personalized flow simulations

Amanda Randles
Biomedical Engineering
Duke University

Why should you care about personalized flow models?

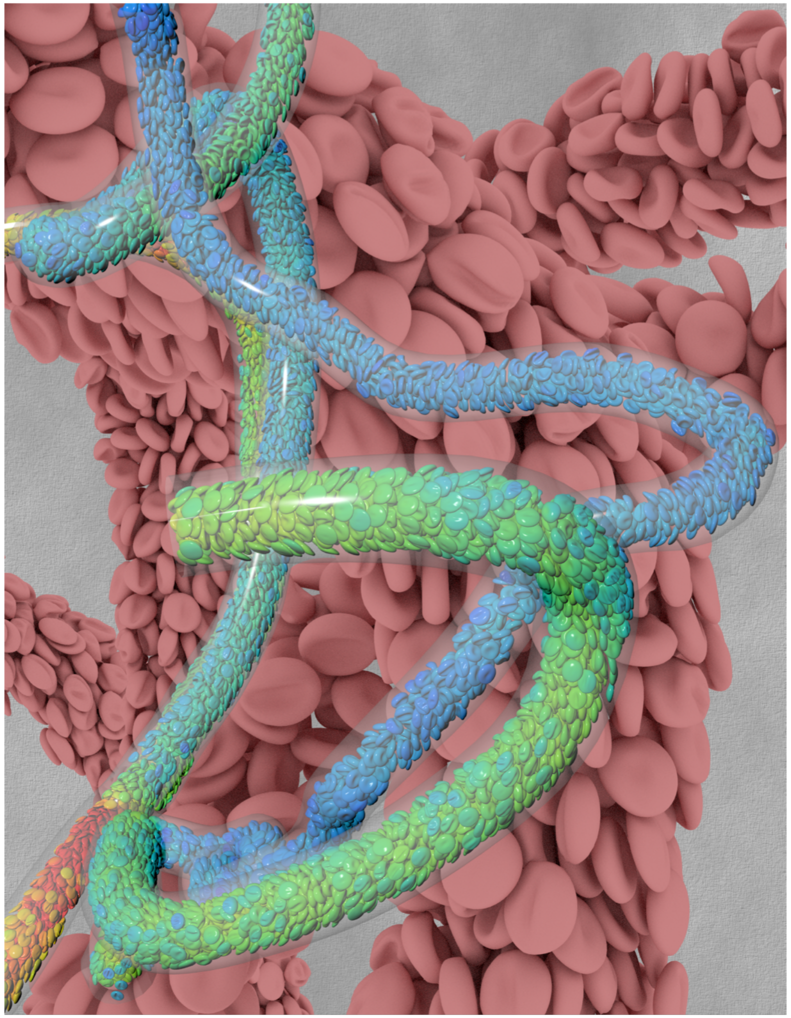

Me in one slide

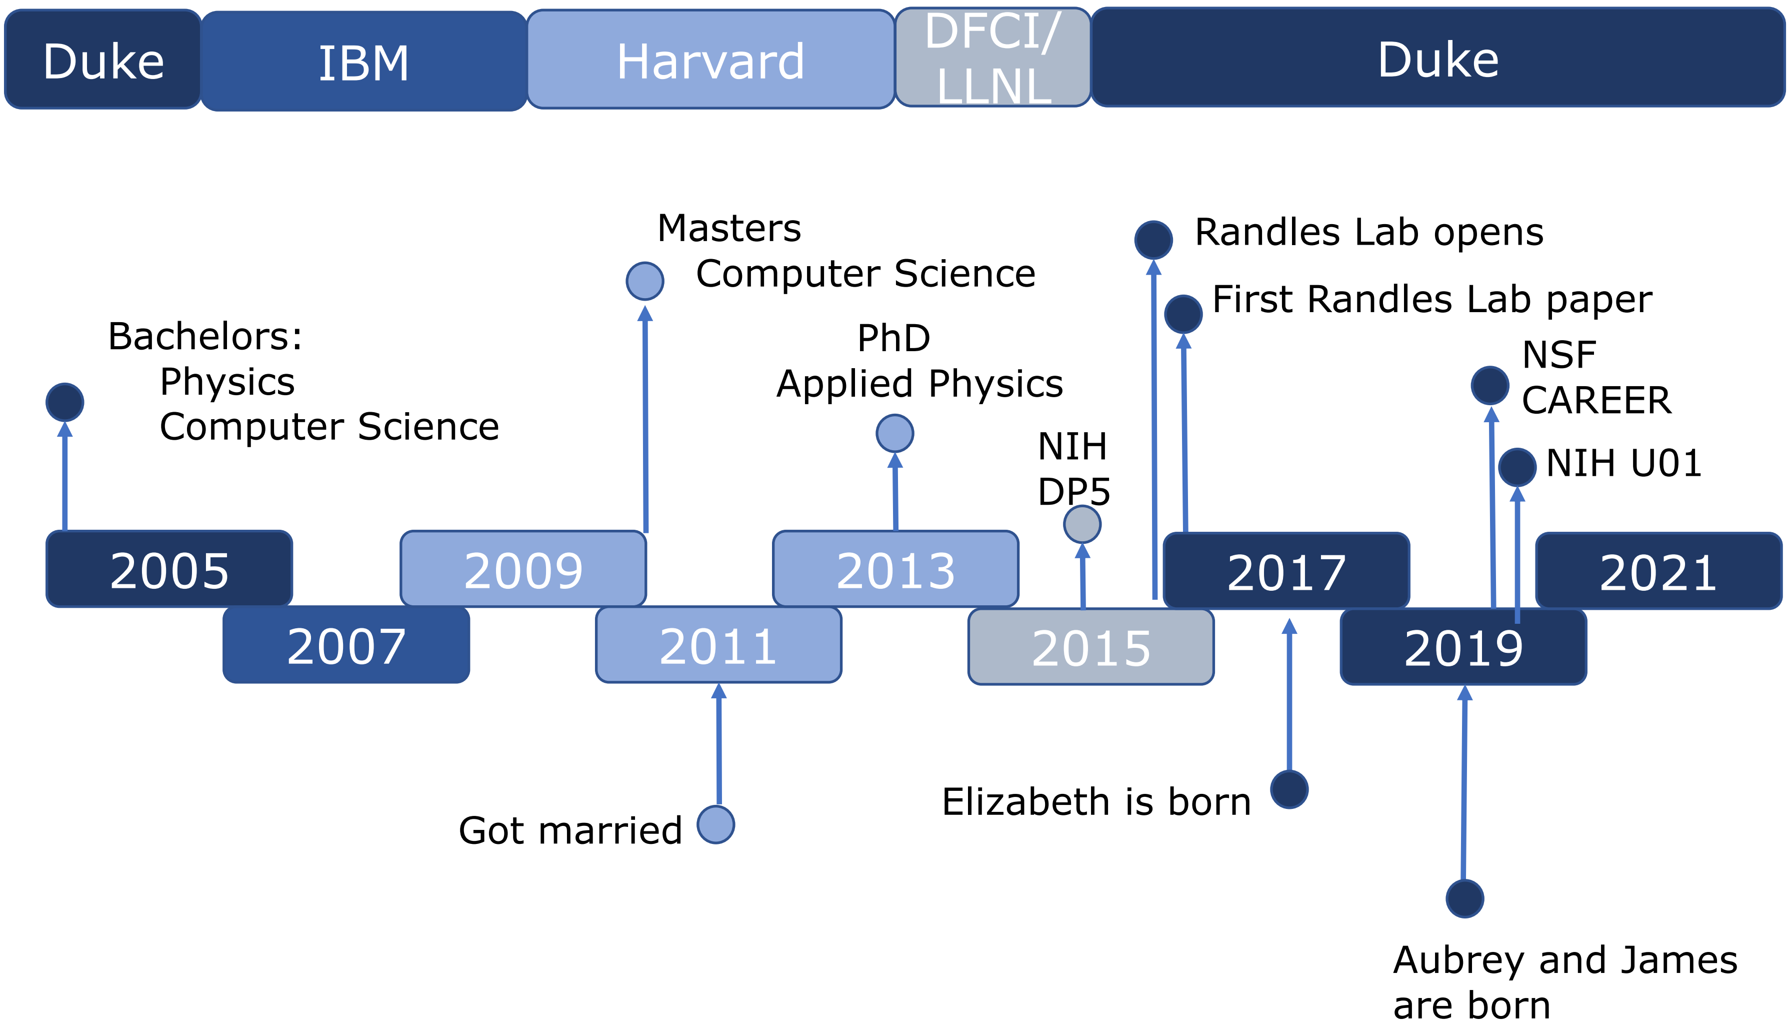

PATIENT-SPECIFIC COMPUTATIONAL MODELS

Patient-derived imaging data

Data Segmentation

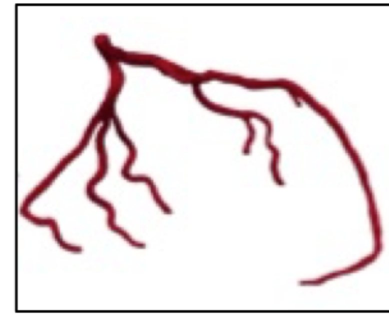

Patient-specific 3D geometries

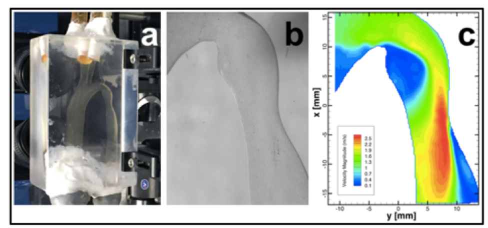

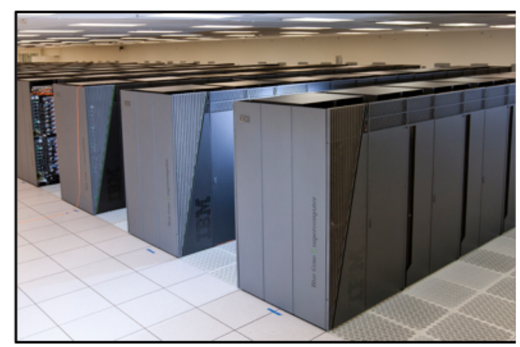

Feedback

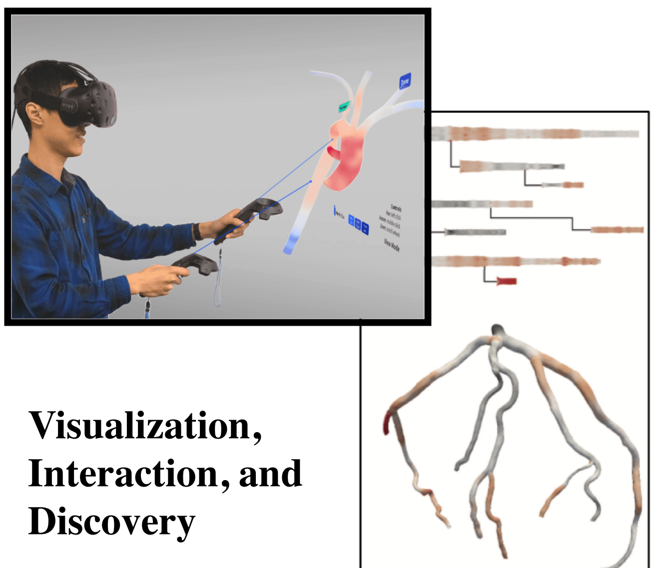

Visualization, Interaction, and Discovery

PATIENT-SPECIFIC COMPUTATIONAL MODELS

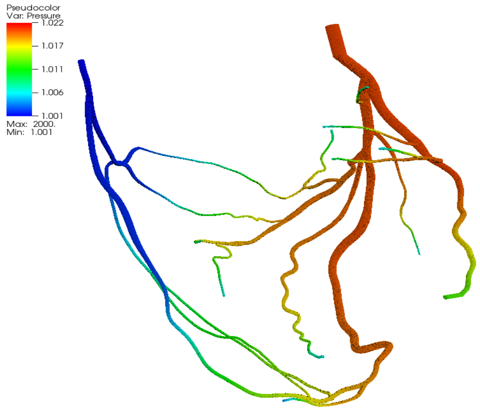

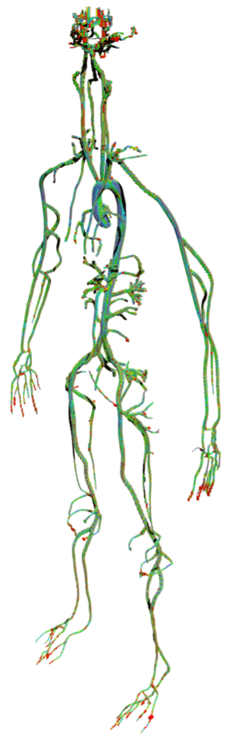

Diagnostics

Treatment Planning

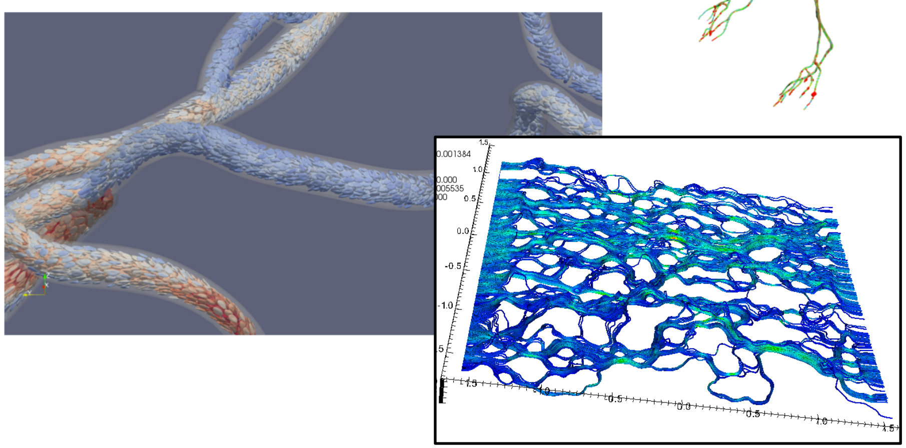

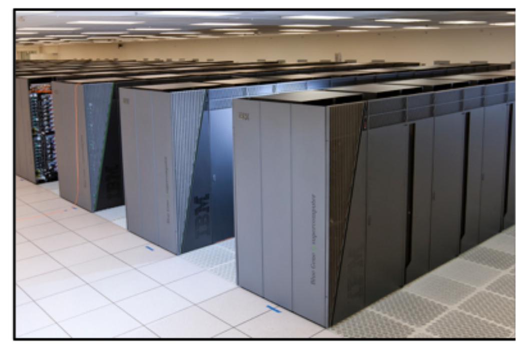

Mechanistic

Massively Parallel Computing

OUTLINE

- Impact of increase in compute capabilities
- Method overview
- Application vignettes:
  - Vascular diseases
  - Fluid structure interaction
  - Ventilator splitting model

IMPACT OF INCREASE COMPUTE POWER

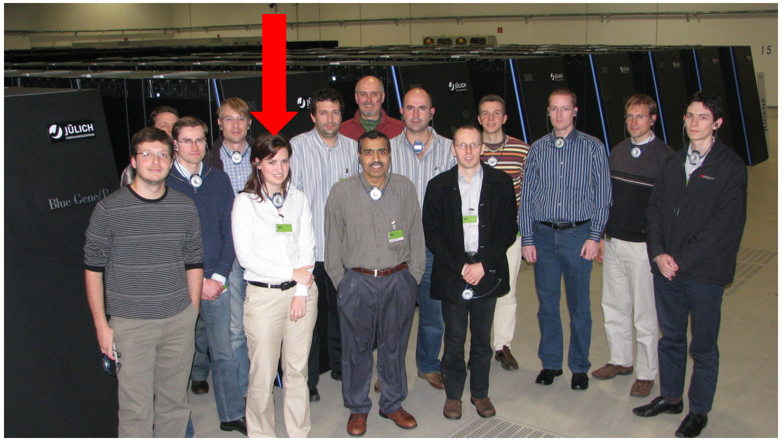

- 2005
  - 1 processor
- 2005-2008
  - Worked at IBM on the Blue Gene supercomputer
  - 65k processors
- 2009 and 2010
  - Juelich Extreme Scaling Workshops
  - ~300k processors
- 2015
  - Worked at LLNL
  - ~1.6 million processors
- 2021
  - Aurora Early Science
  - Exascale

SIMULATIONS ARE COMPUTATIONALLY INTENSE

FULL ARTERIAL NETWORK

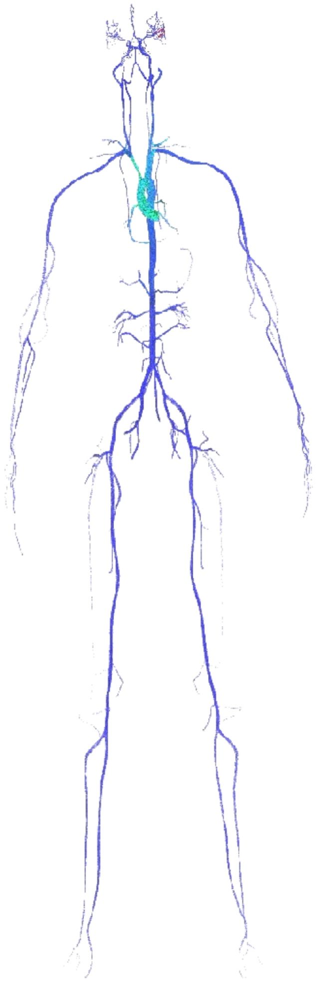

COMPUTER SCIENCE CHALLENGES

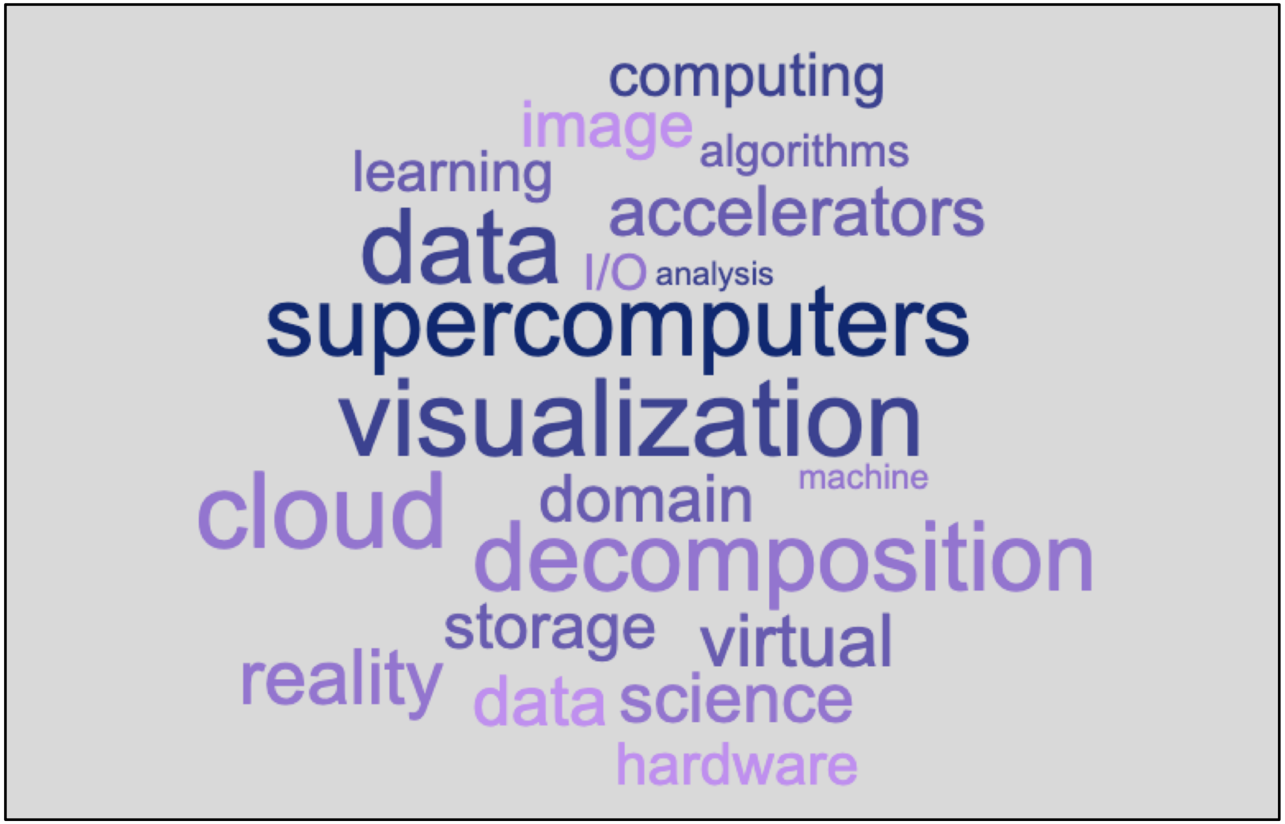

Computational Fluid Dynamics

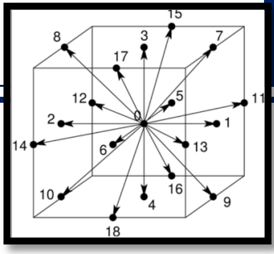

- Solves the (weakly compressible) Navier-Stokes equations
- Minimal communication between lattice points during update
- Macroscopic quantities computed at lattice points

$$\rho = \sum_i f_i \qquad \vec{u} = \frac{1}{\rho} \sum_i \vec{c}_i f_i \qquad P = c_s^2 \rho$$

- Straight forward treatment of complex geometry

How do we make these models tractable?

DOMAIN DECOMPOSITON

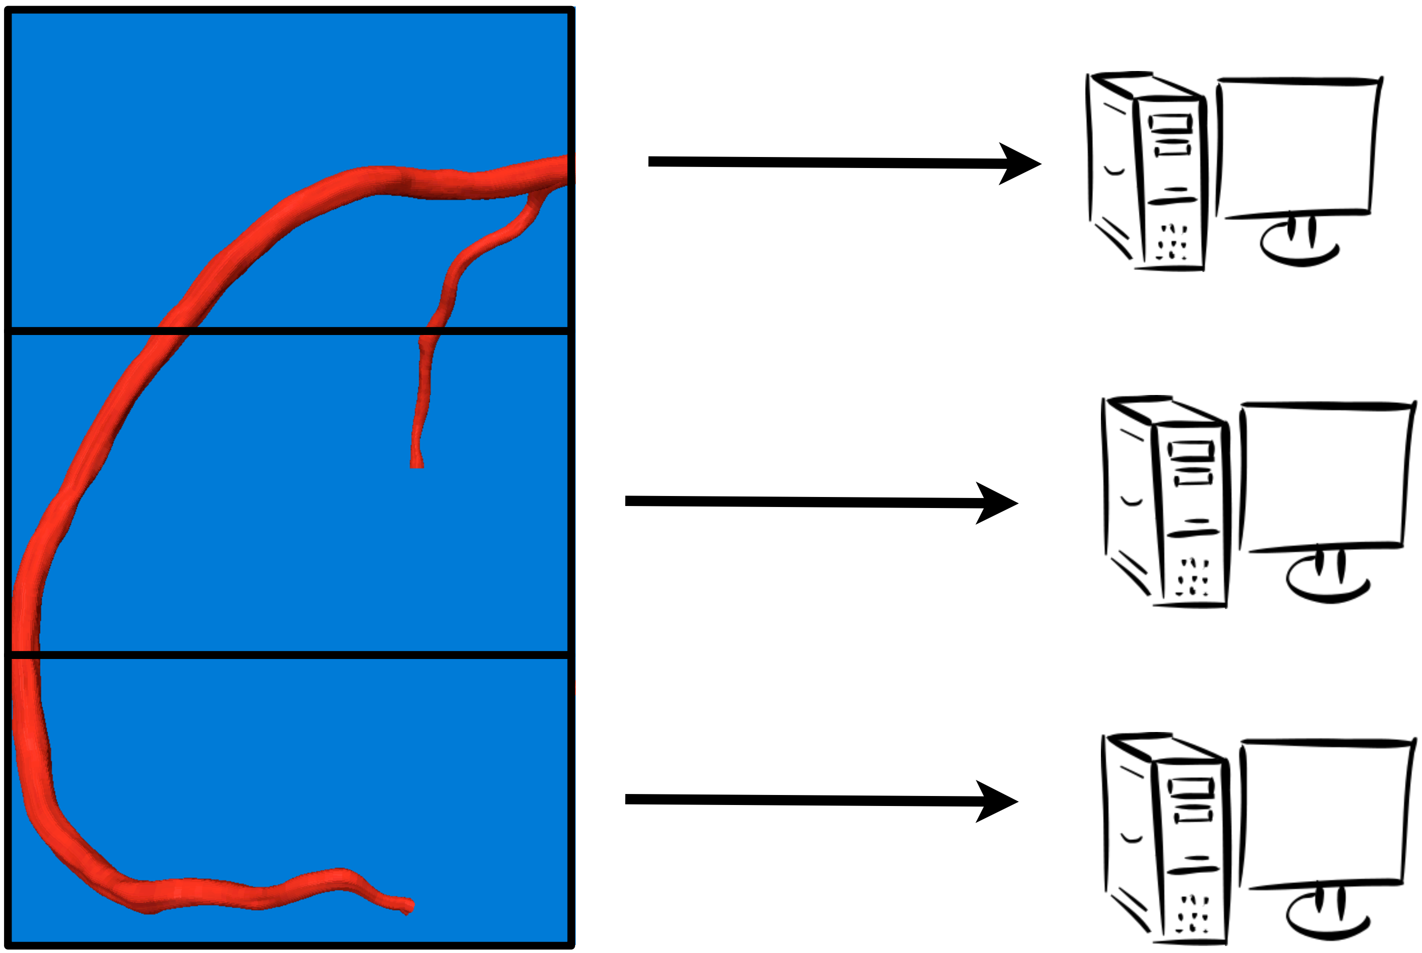

IRREGULAR DOMAIN DECOMPOSITION

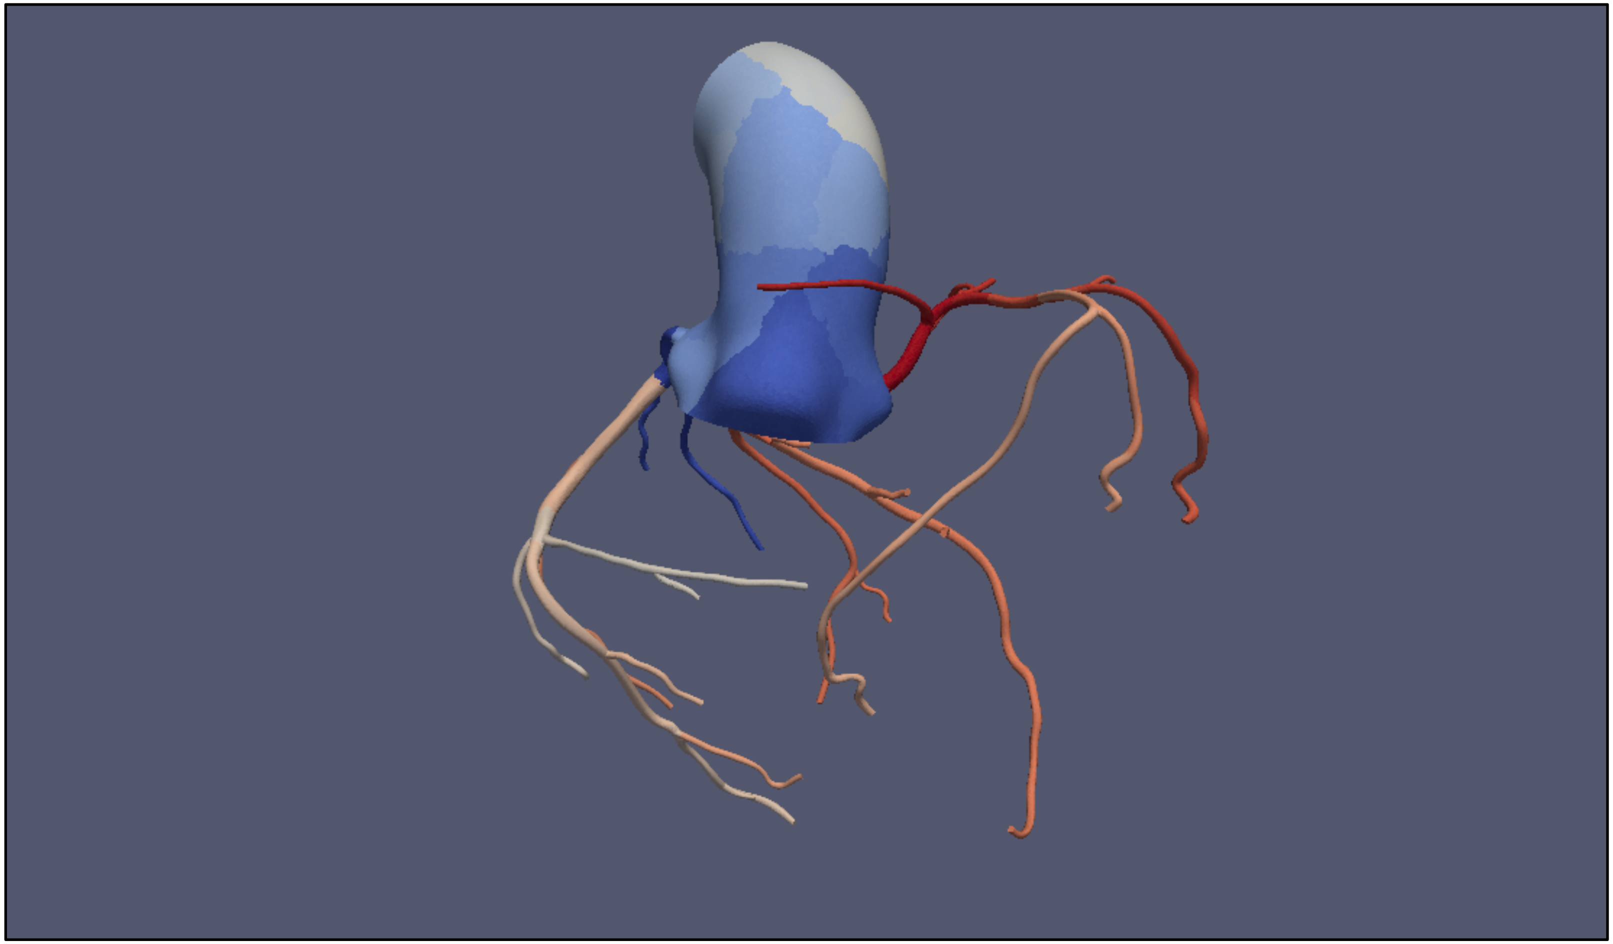

MINIMIZING MEMORY FOOTPRINT

Full body arterial	20 micron	9 micron
Data Grid		68909 x 25107 x 188584
Data memory		90.2 PB
Fluid nodes		509 billion
Fluid memory		140.7 TB
Fluid Fraction		0.15%

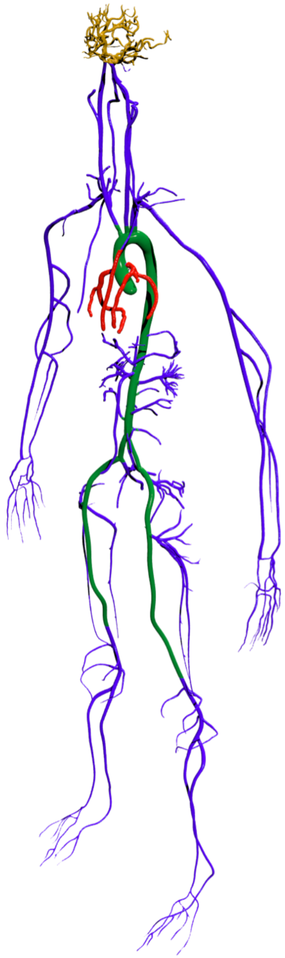

- Sequoia Blue Gene/Q total system memory: 1.6 PB
- Indirect addressing is mandatory, and only the first step
- Initialization and load balance now significantly more challenging

MEMORY LAYOUT SCHEMES

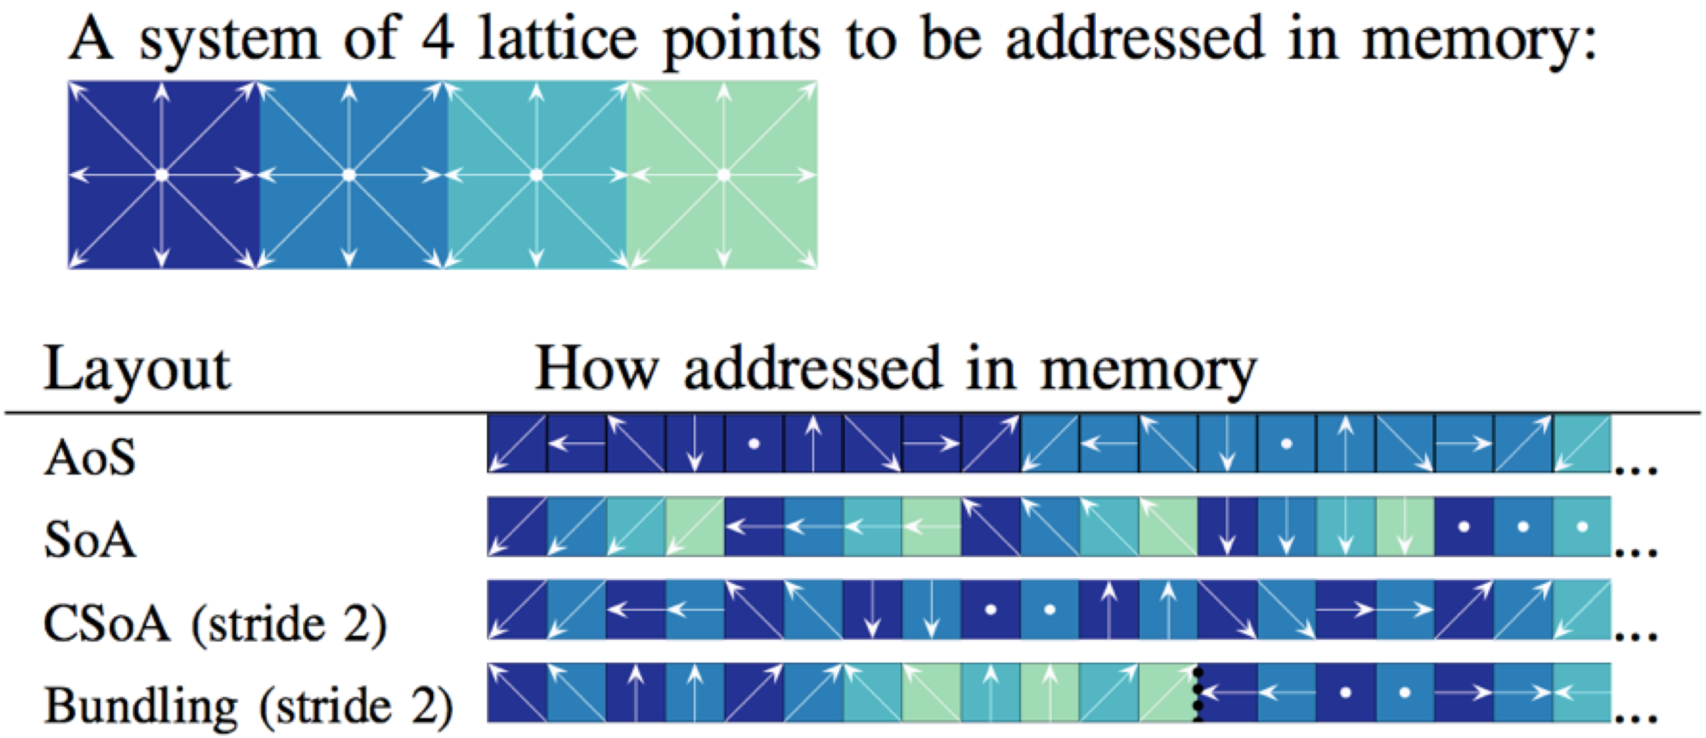

Herschlag *et al.* IPDPS 2018

Efficient Use of Large-Scale Supercomputers

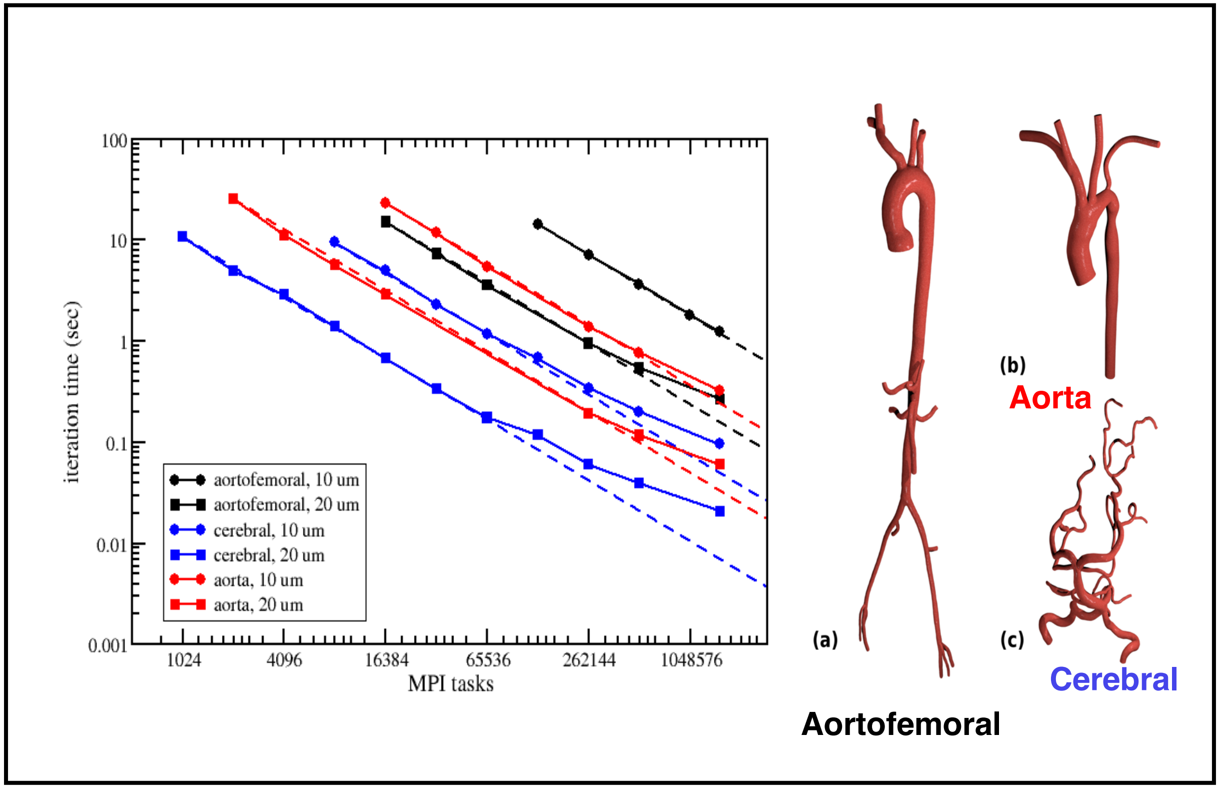

Randles *et al*. Journal of Computational Science, 2015

Google Cloud

Azure

aws

APPLICATIONS OF MASSIVELY PARALLEL FLUID MODELS

Cardiovascular Disease

CADs share highest burden of heart diseases with 370,000 death per year in the United States.

Personalized simulations can provide lesion-specific data to determine the functional severity of a stenosis and guide therapeutic action.

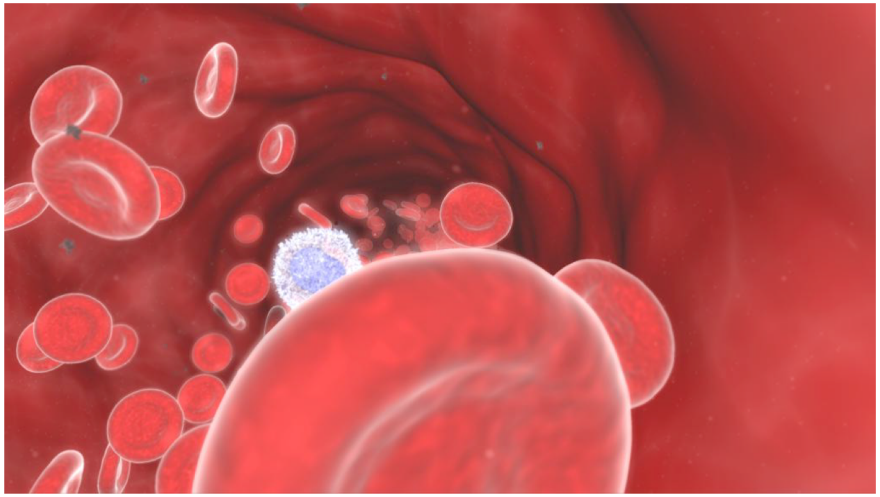

50% of men and 64% of women who die of cardiovascular disease have no previously detected symptoms

Diagnostics

Determining lesion ischema: to stent or not?

$$FFR = \frac{\text{Distal Coronary Pressure (Pd)}}{\text{Proximal Coronary Pressure (Pa)}}$$

(During Maximum Hyperemia)

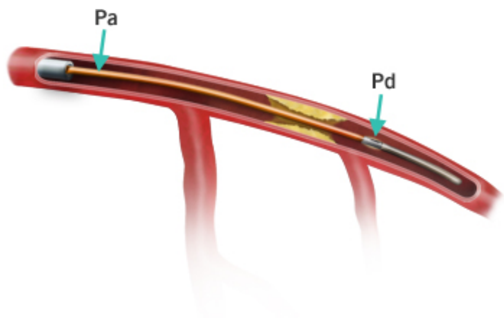

1. COURAGE NEJM 2007
2. DEFER Study JACC 2007
3. FAME NEJM 2009
4. FAME II NEJM 2012
5. FAME III AHJ 2015
6. Clinical outcomes FAME II, Circ. 2019

A

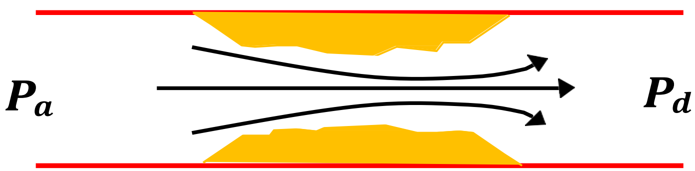

Hyperemia: Large gradient

B

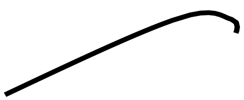

Pressure guide-wire

+

Hyperemic drug

=

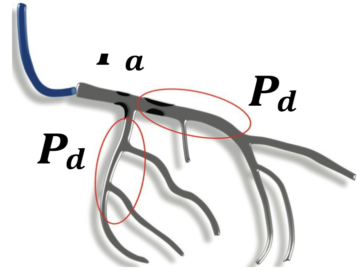

$$Invasive\ \boldsymbol{FFR} = \frac{P_d}{P_a}$$

FFR is the current gold standard for the assessment of coronary artery disease

Deriving flow from angiography data

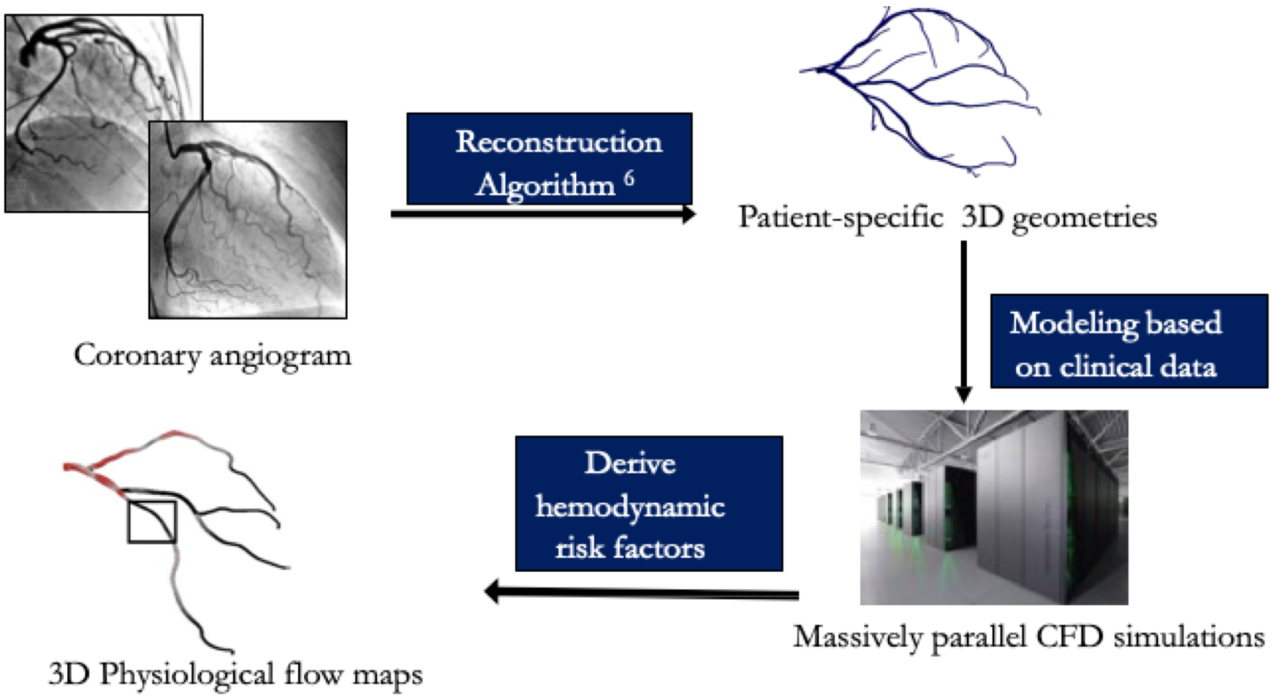

Significant physiology can be resolved using CA applied to CFD, which can help determine intracoronary hemodynamics more accurately and augment FFR

Vardhan et al. Nature Scientific Reports 2019

Importance of including side branches

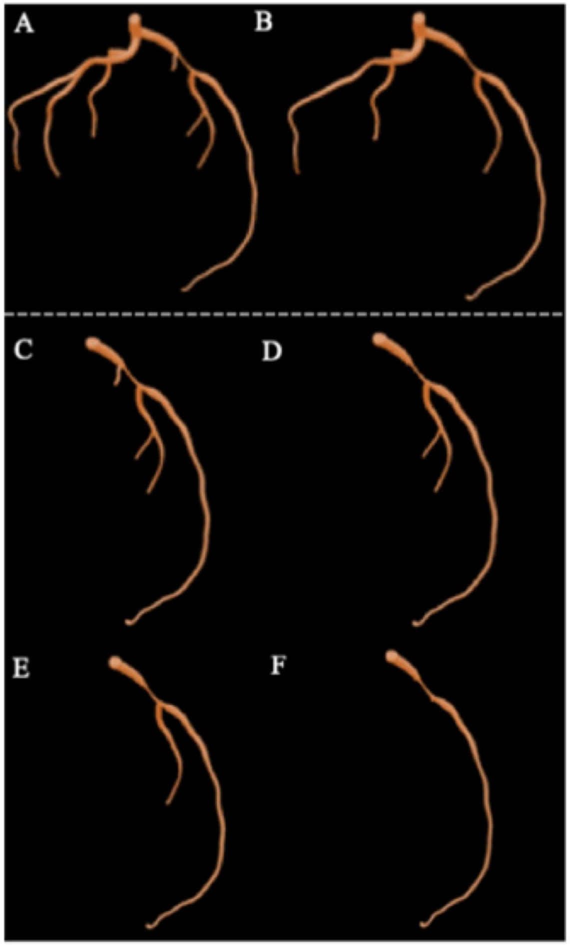

Vardhan et al. Nature Scientific Reports 2019

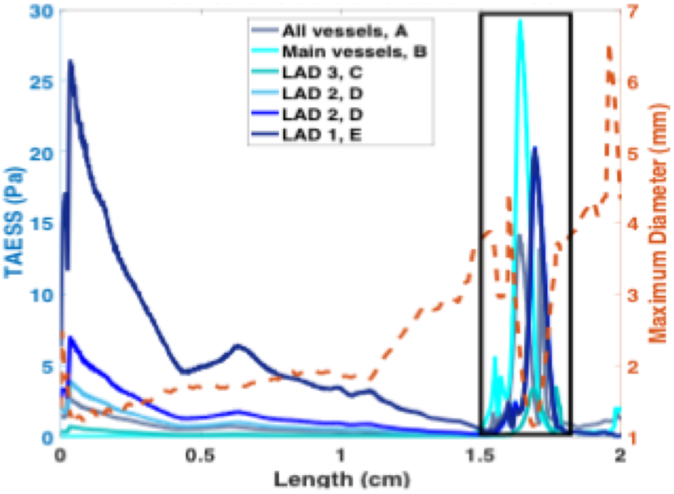

Capturing the full arterial tree including side branches is critical to accurately assessing flow characteristics.

Capturing complex lesions

- Complex lesions found in more than 20% of patients
- High risk of developing secondary adverse cardiac events
- Commonly excluded from most non-invasive clinical trials

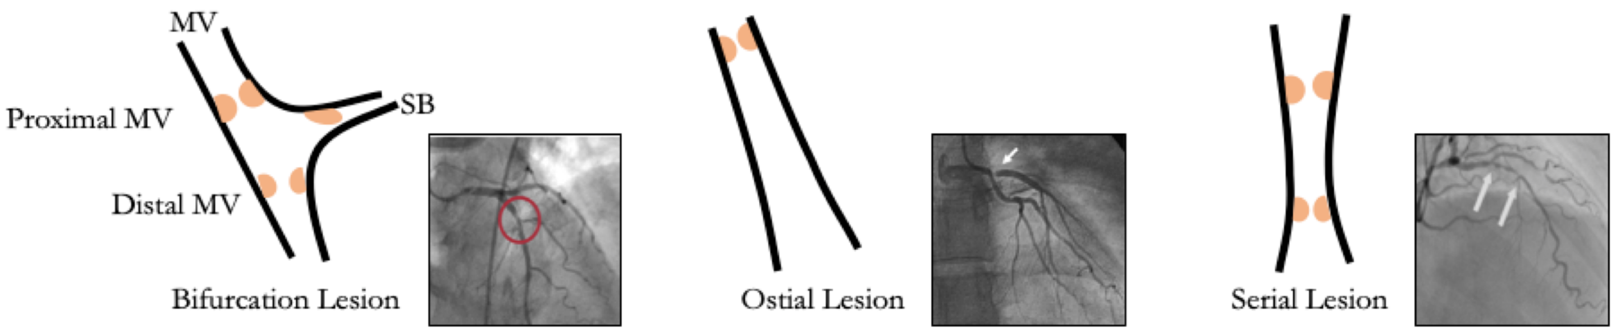

Determining hemodynamic precarity in complex lesions can help understand disease progression and guide future treatment decisions

Vardhan et al. Nat. Sci. Rep. (2021)

Matching CA-CFD to *in vivo* data

Case	Vessel	Clinical Resting Pressure	CFD Resting Pressure	Percent Error
Case 1	LAD	0.93	0.94	0.6%
Case 2	RCA	0.97	0.98	1.0%
Case 3	LCx	0.96	0.97	1.0%
Case 4	RCA	0.88	0.88	0.0%
Case 5	LAD	0.95	0.95	0.0%
Case 6	RCA	0.91	0.90	0.7%
Case 7	Left Main	0.96	0.99	3.1%
Case 8	LCx	1	0.98	2.0%
Case 9	RCA	0.98	0.92	6.0%
Case 10	LAD	0.82	0.72	12.2%
Case 11	LAD	0.92	0.95	3.3%
Case 12	RCA	0.98	0.93	5.1%
Case 13	LAD	0.93	0.97	4.3%
Case 14	RCA	1	0.95	5.0%
				3.16%

Invasive resting gradient and CA-CFD resting gradient are in close agreement with each other

Vardhan *et al*. Nature Scientific Reports 2021

Mixed order models: VA-ECMO

VA-ECMO – mechanical support system

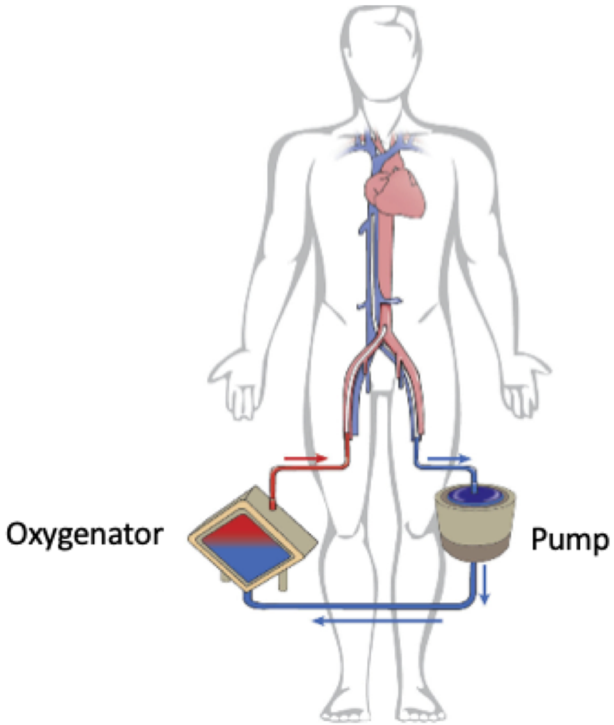

- External heart and/or lungs for patients with cardiopulmonary failure
- Process
  1. Drain deoxygenated blood from veins
  2. Pump
  3. Oxygenate
  4. Infuse blood through insertion cannula in femoral artery, axillary artery, or aorta

Importance of the mixing zone

- ~15% patients develop neurological complications
- Cerebral hypoxia
  - Differential hypoxia - well oxygenated VA-ECMO flow meets poorly oxygenated cardiac output
  - Mixing zone – region of mixing in the aorta

Important to determine how physician tunable parameters and cannulation location influence flow to the brain and location of mixing zone

Building 1D-3D coupled models

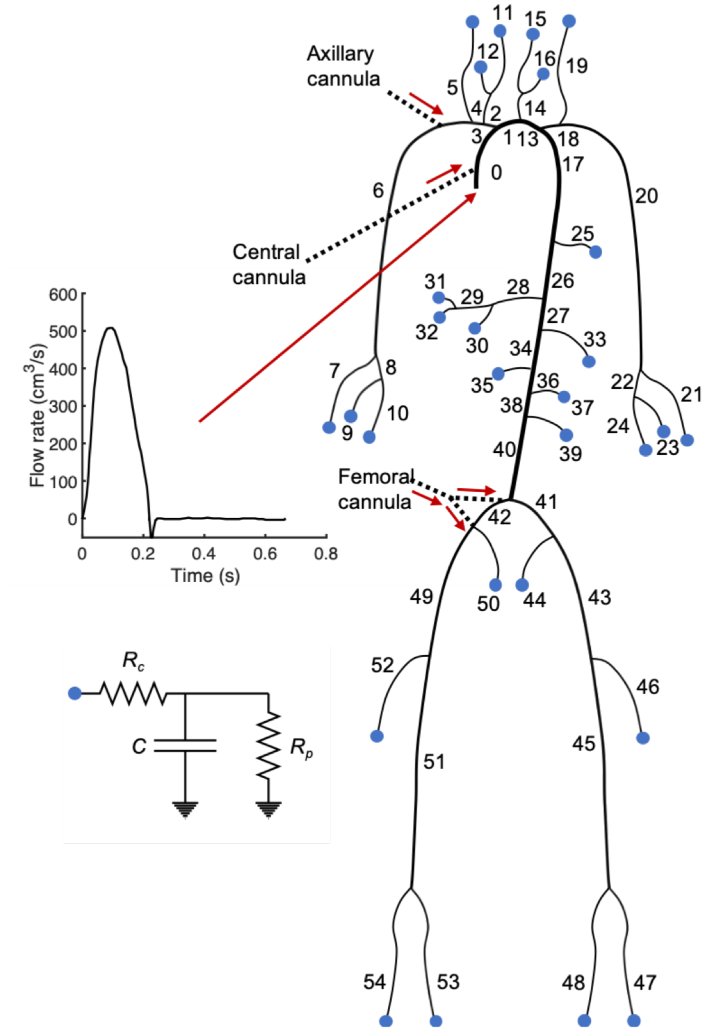

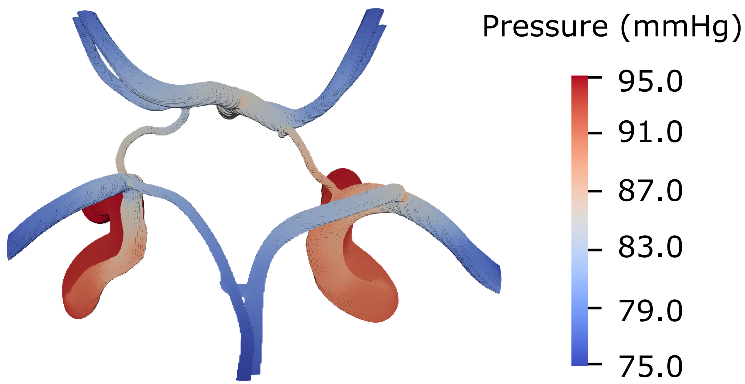

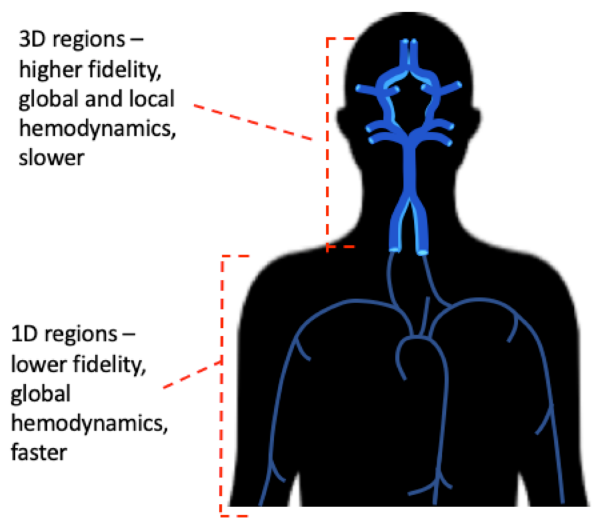

Feiger *et al.* J Biomech 2020
Feiger *et al.* Comp in Bio and Med. 2020

Building 1D-3D coupled models

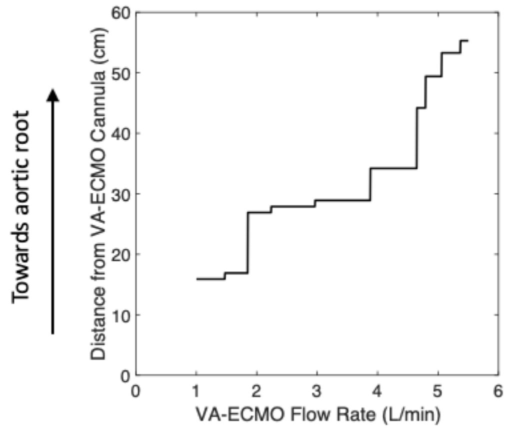

- Coupled models provide accurate 3D hemodynamics and capture VA-ECMO properties
- VA-ECMO flow rate drives mixing zone
- High flow rates are needed to oxygenate the brain

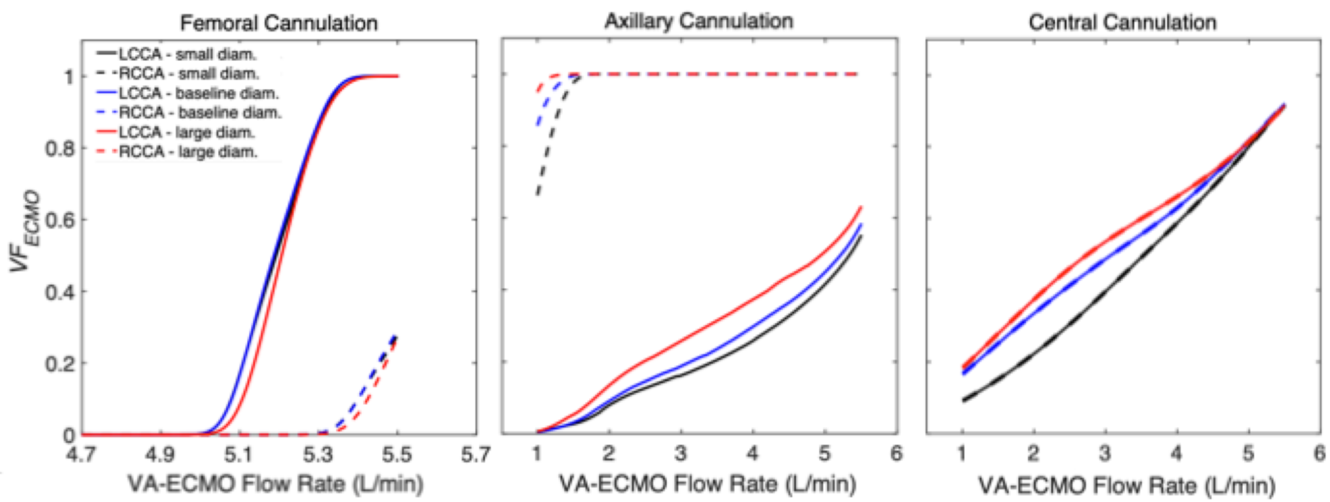

Feiger *et al.* J Biomech 2020
Feiger *et al.* Comp in Bio and Med. 2020

Treatment Planning

Results from CFD simulations can guide clinical treatment.

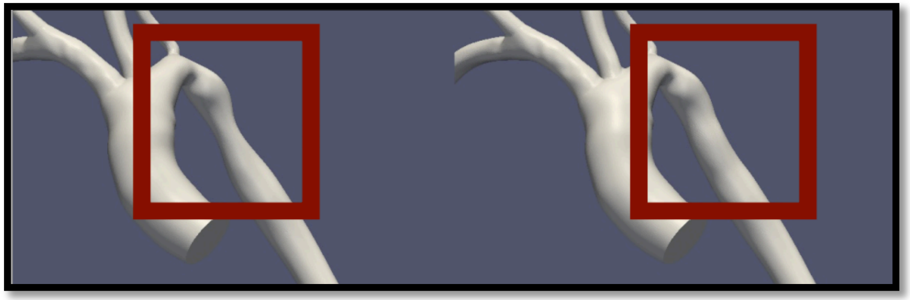

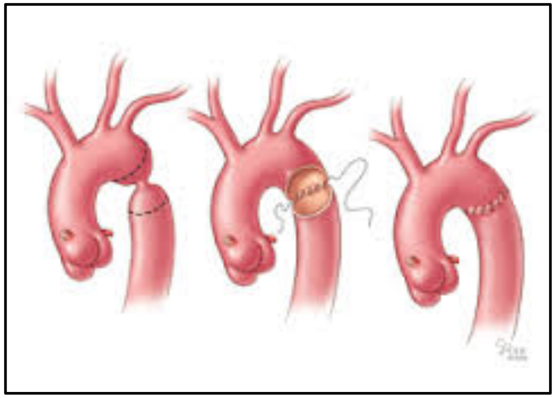

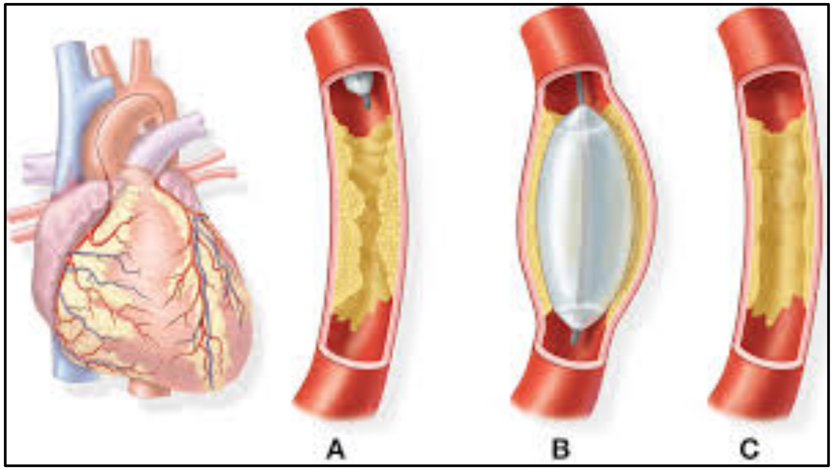

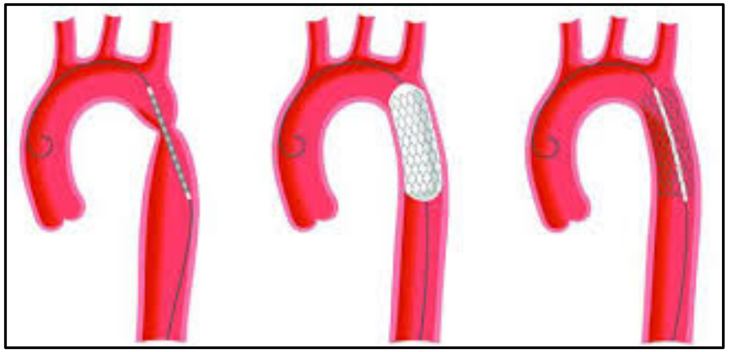

Simulations can capture local effects on ABI

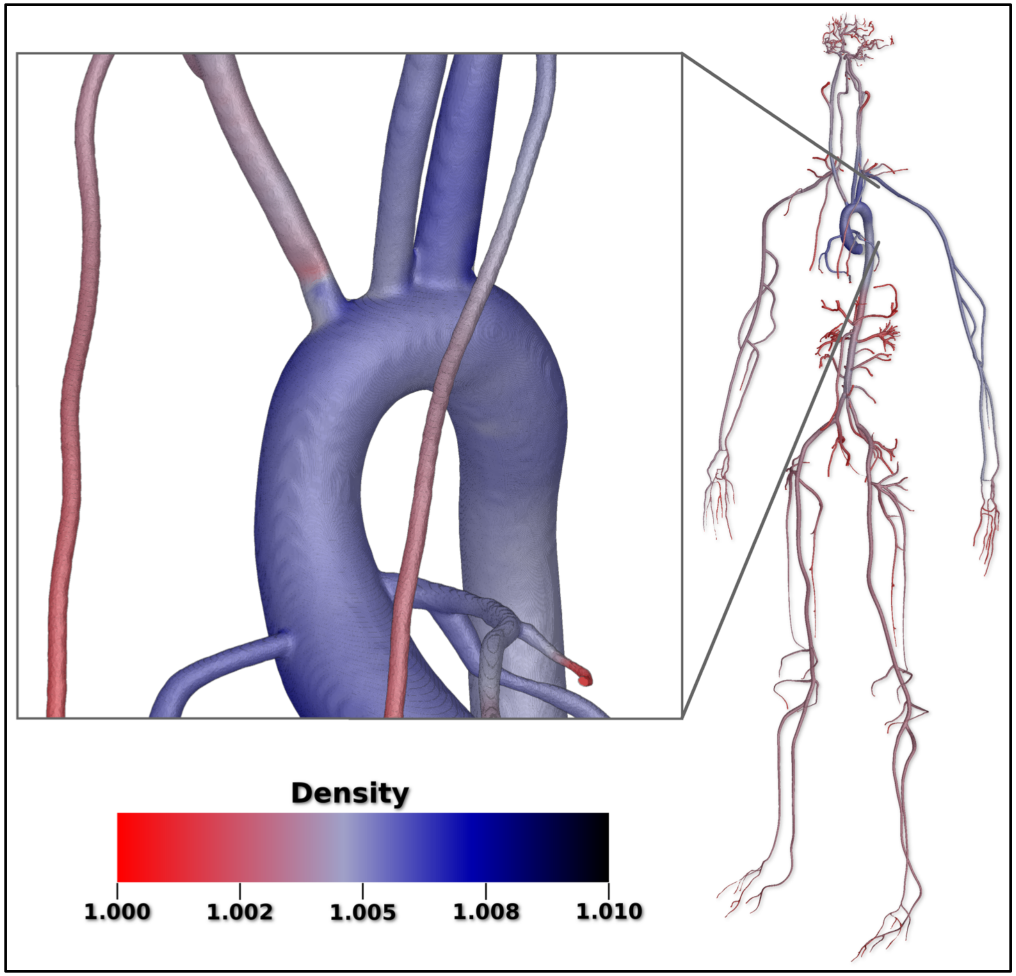

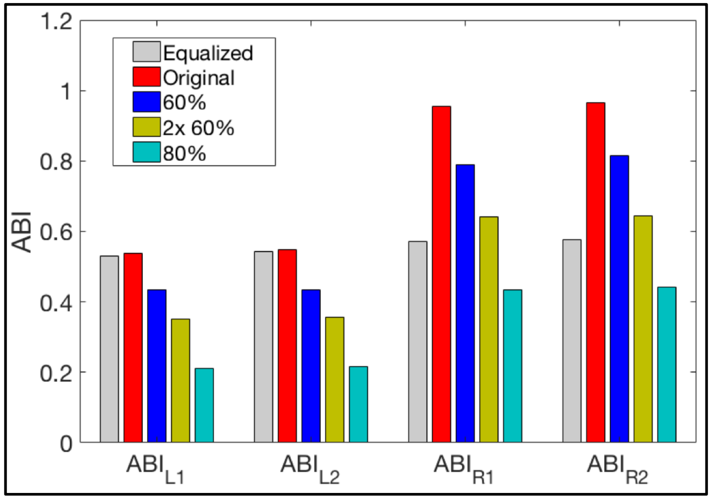

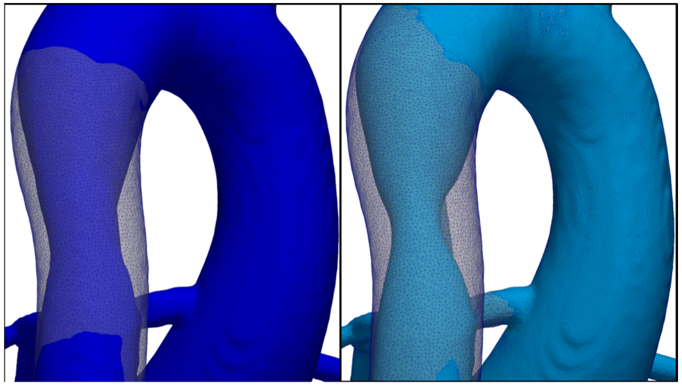

Gounley et al. Journal of Biomechanics 2019

HarVis: XR interaction

- Using devices like the Occulus Rift, HTC Vive, and zSpace to investigate different modes of interaction

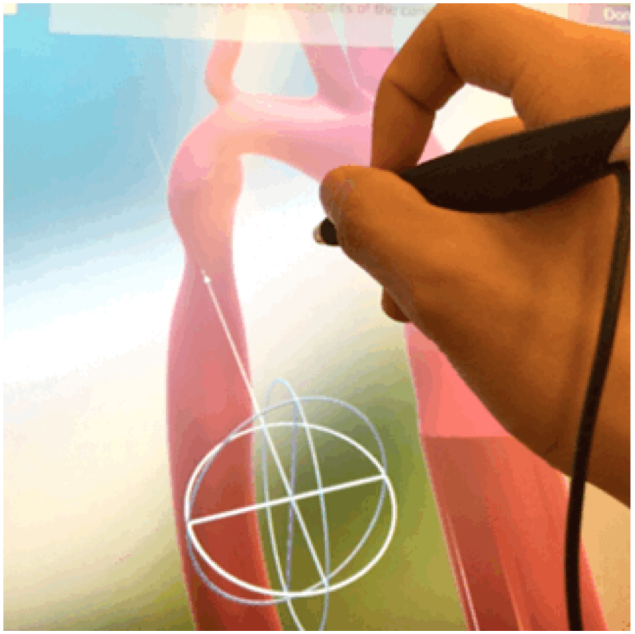

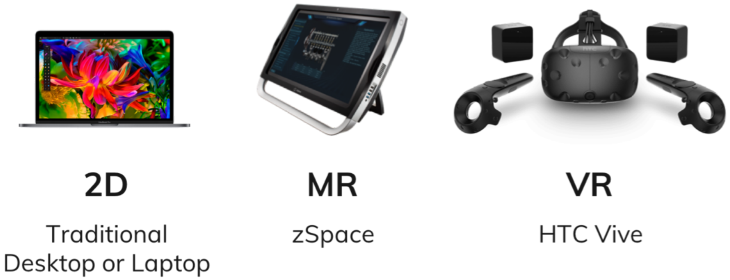

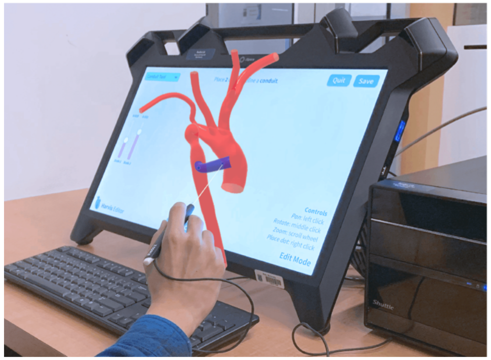

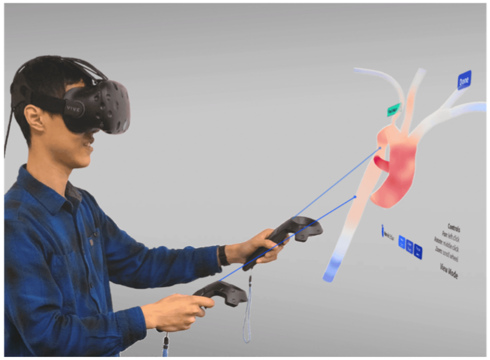

Vardhan *et al.* MICCAI 2019
Shi *et al*. Journal of Computational Science 2020

Assessing user interaction

- Full immersive devices can provide reduced error for completing defined tasks like conduit placement or virtual revascularization
- Fully immersive devices prove most effective for PCI planning
- Access to shear stress maps reduces differences in participant accuracy across devices

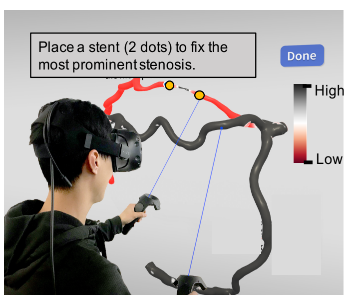

Fully Immersive
With WSS

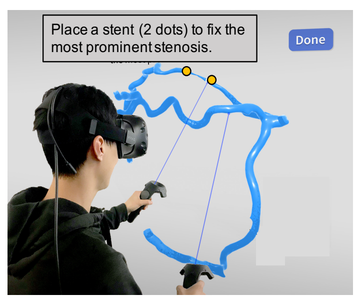

Fully Immersive
No WSS

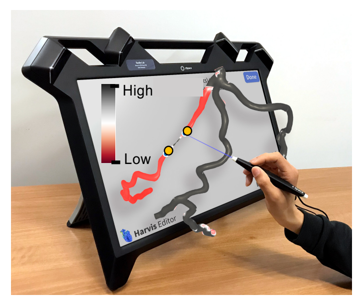

Semi Immersive
With WSS

Semi Immersive
No WSS

Vardhan, *et al.* MICCAI 2019
Shi *et al*. Journal of Computational Science 2020

Holistic view of the patient

Coarctation of the aorta

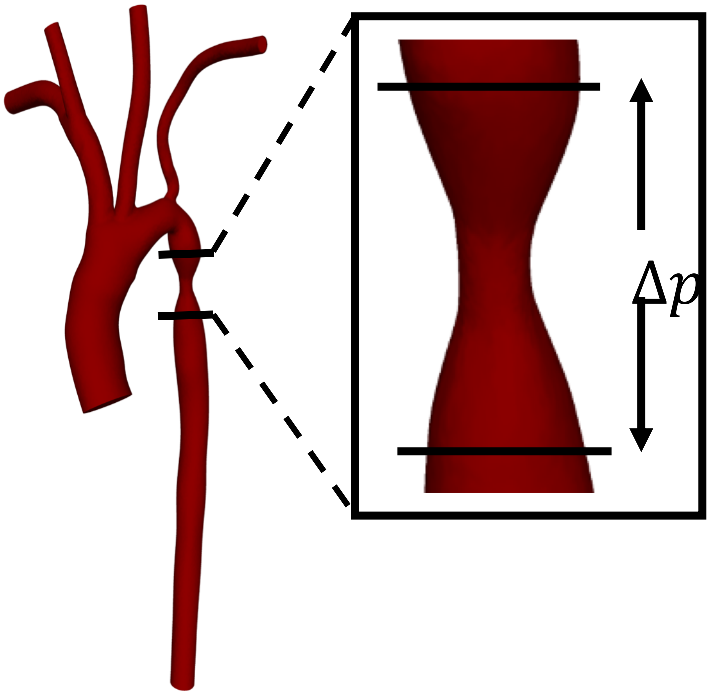

- Congenital heart defect
- Stenosis in descending aorta
- Affects 3,000 – 5,000 patients in U.S. each year
- Intervention if **Δp > 20 mmHg**
- TAWSS –> plaque progression

Machine learning pipeline

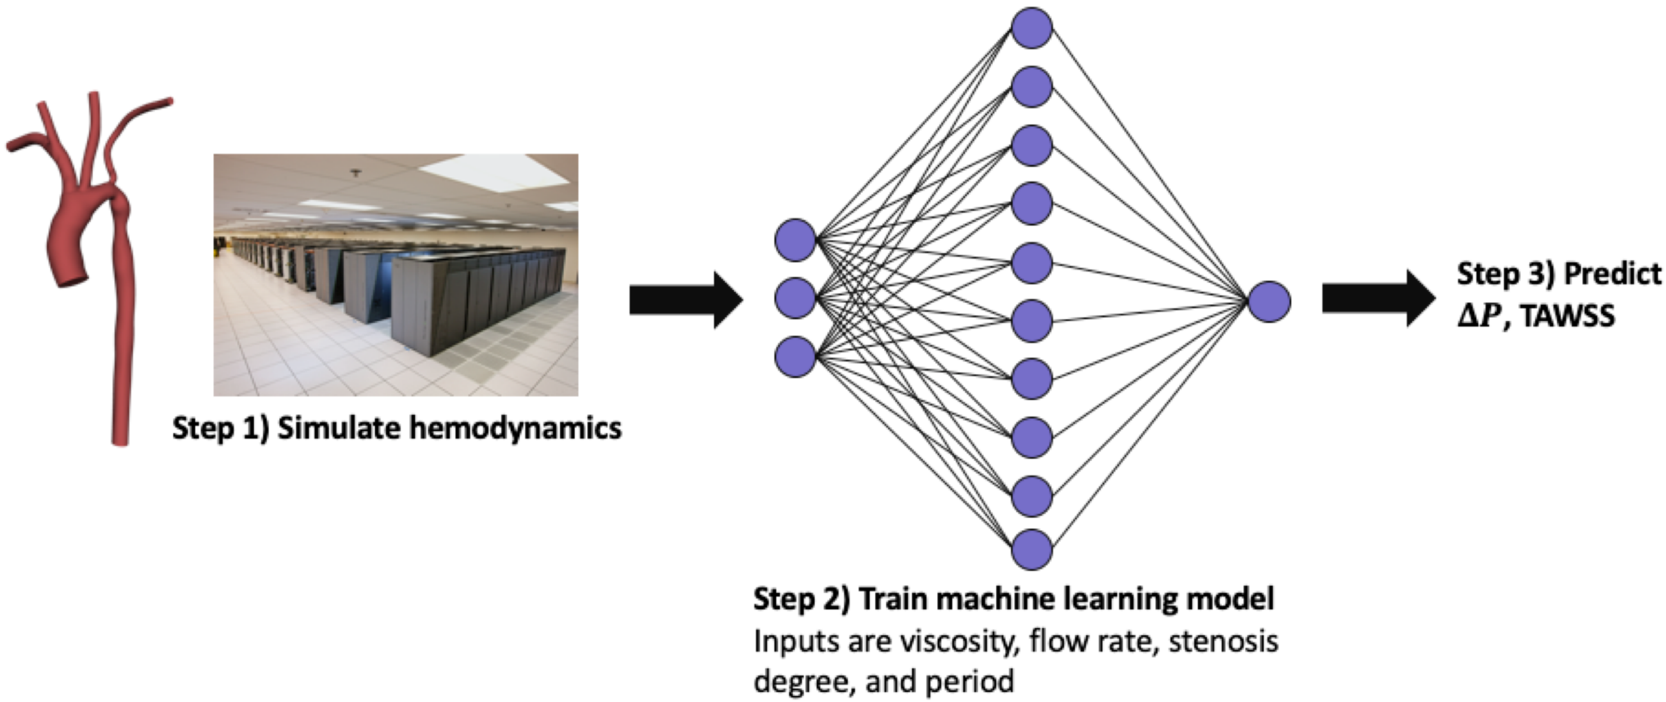

Feiger *et al*. Nature Scientific Reports 2020

AI-driven Simulations: Stratifying Patient Risk

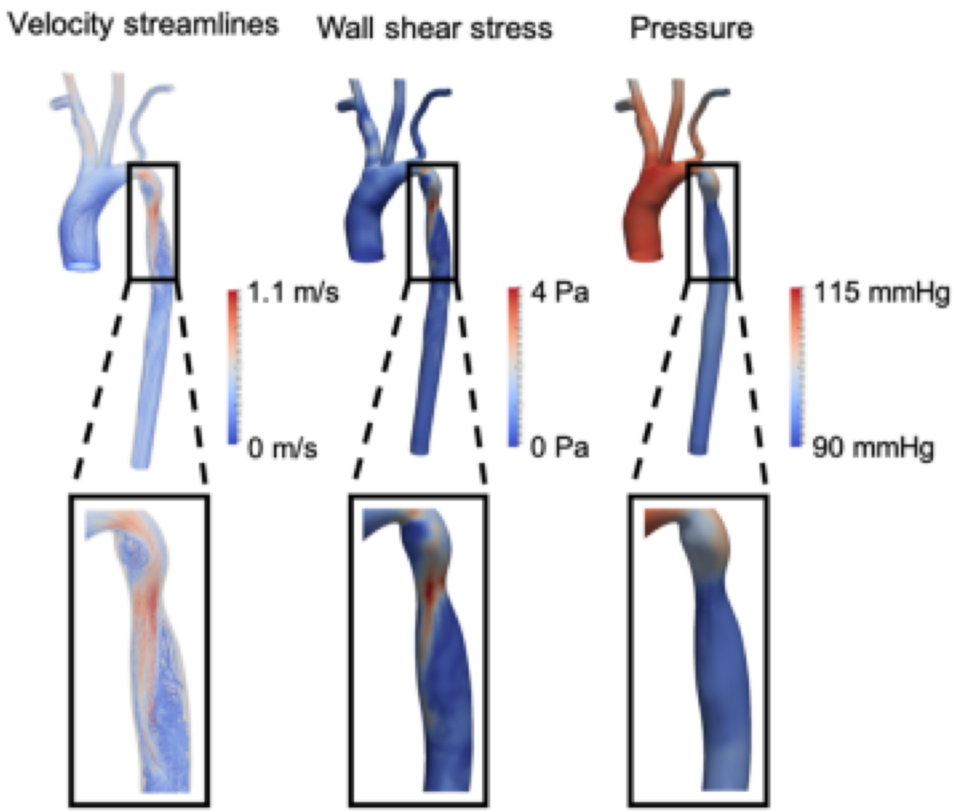

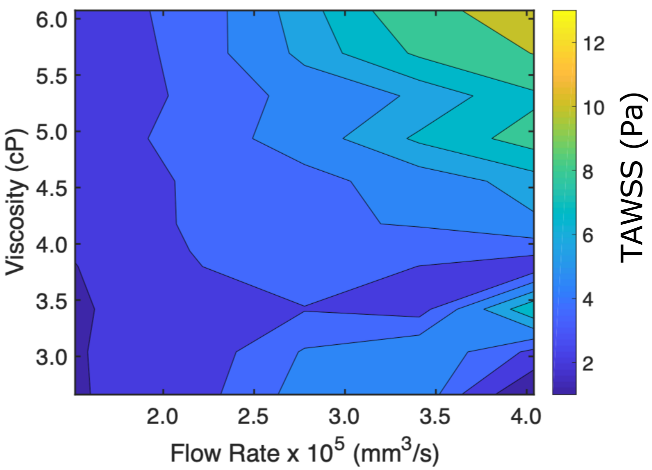

Identified minimal number of simulations needed to stratify patient risk

Feiger *et al*. Nature Scientific Reports 2020

Fluid Structure Interaction Models

Cell-specific models

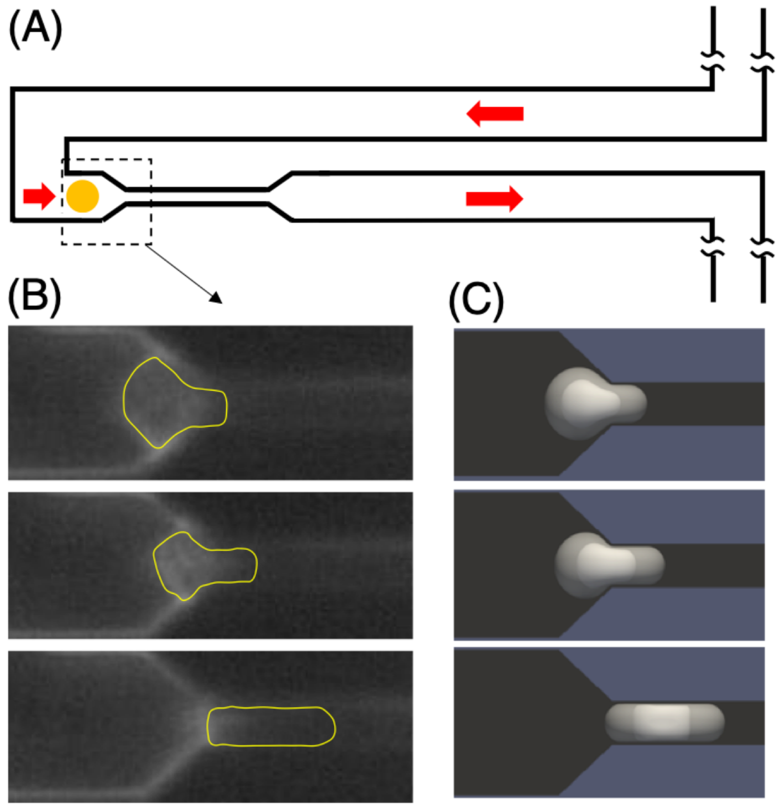

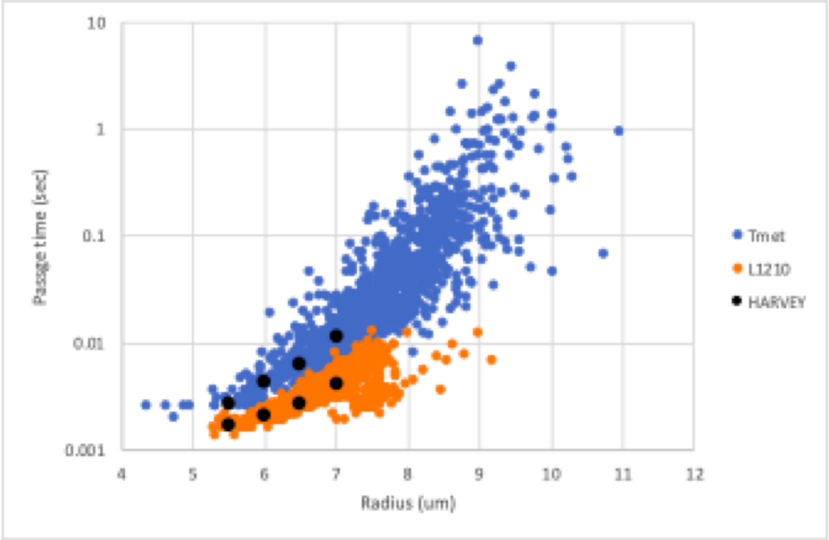

We need a tunable model to capture cancer cell-specific parameters.

Balogh *et al.* Scientific Reports 2021

Predicting cellular attachment

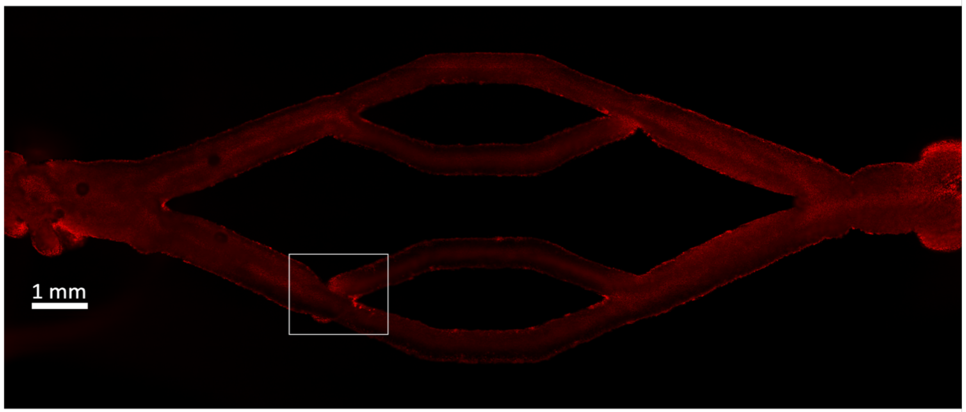

3D Bioprinted Bed

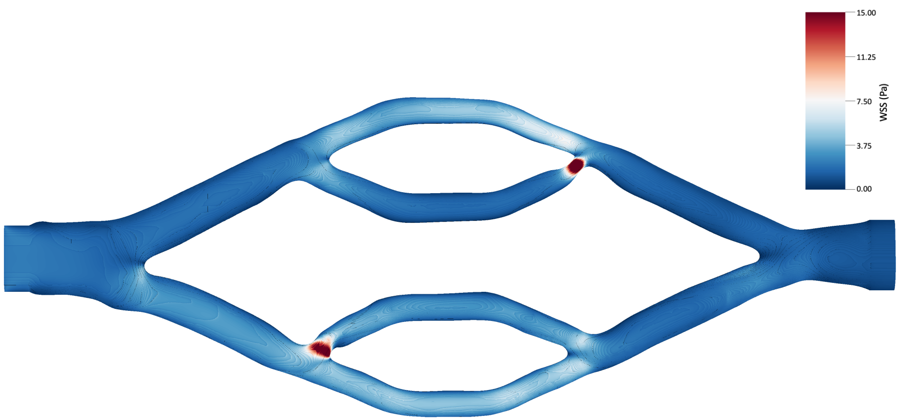

Bulk Fluid Simulation

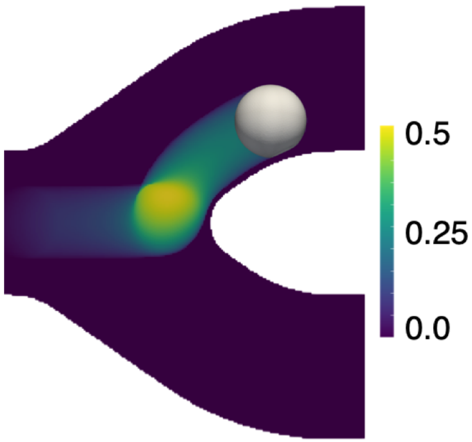

FSI model

Hynes et al. Science Advances 2020
Pepona et al. CMBE 2020

Scaling the LBM-IB model

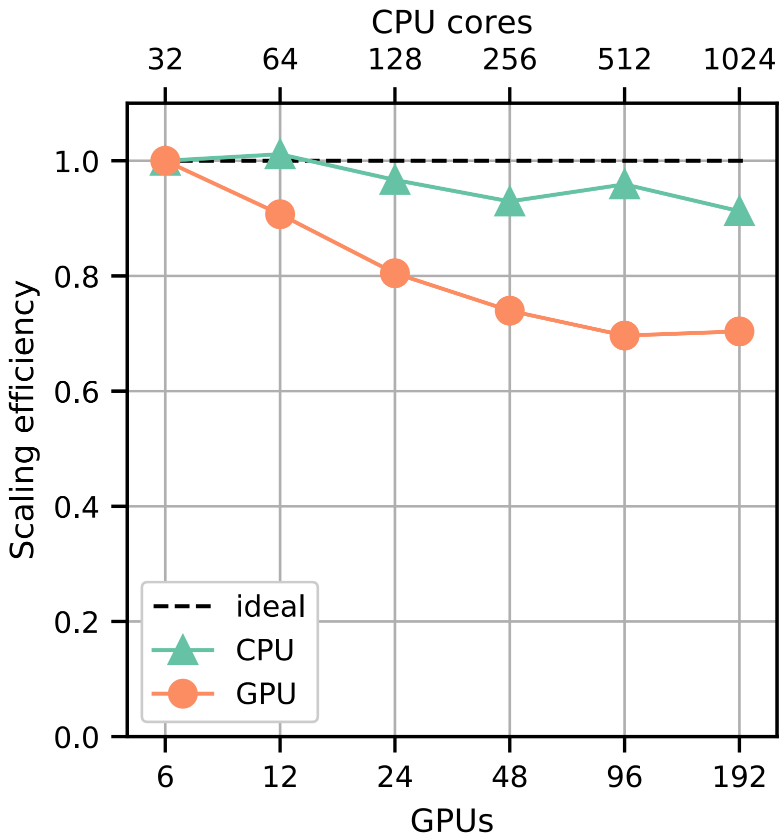

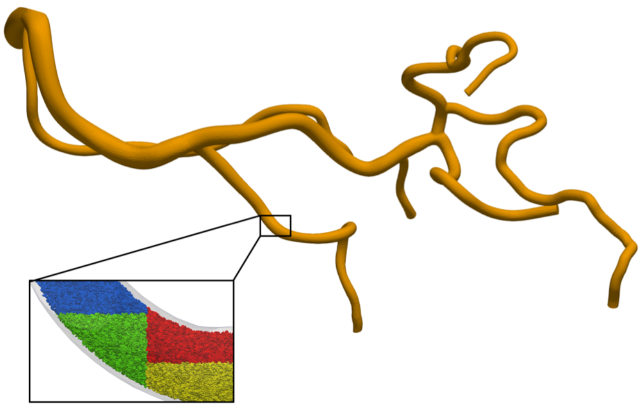

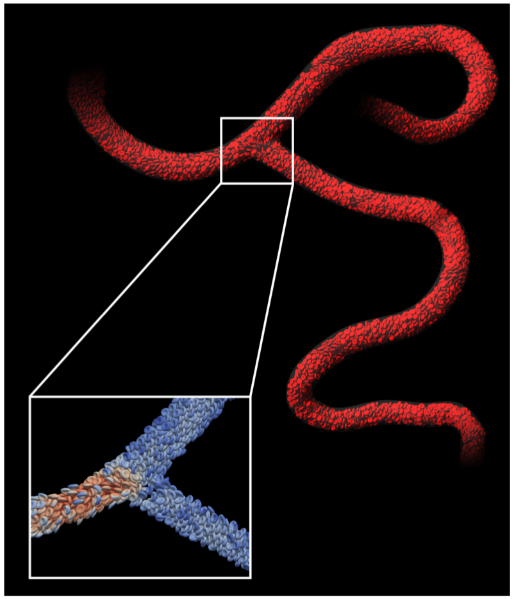

Gounley et al. ICCS 2019
Ames et al. JCS 2020

Adaptive Physics Refinement

- Generate an adaptive domain, multi-physics algorithm to resolve tumor cell trajectories
- Full or partial body simulations in which parts of the arterial system are purely fluid, whereas other parts resolve red blood cells.
- Utilize heterogeneous architectures efficiently.

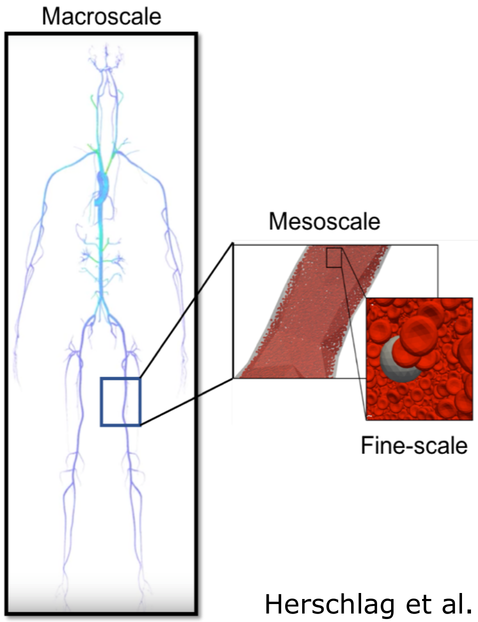

Herschlag et al. IEEE Cluster 2019

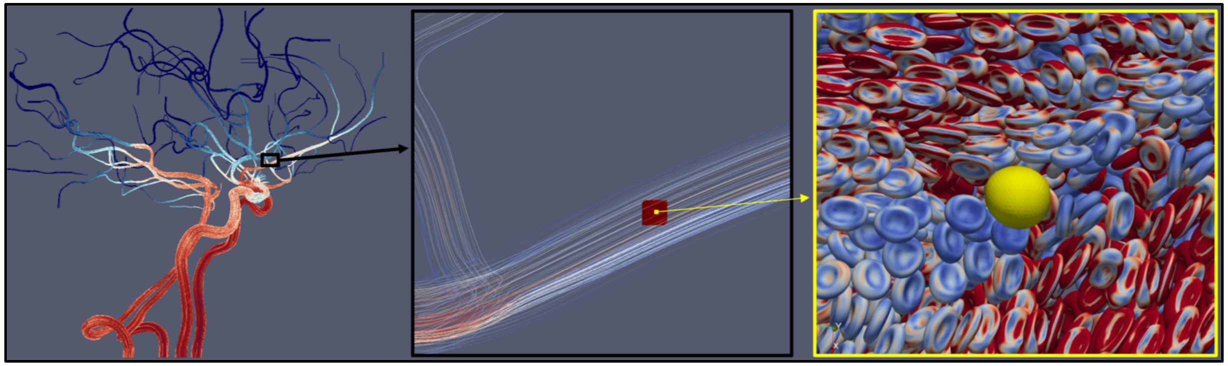

Next steps: preparing for exascale

- Large-scale simulations create petabytes of data per timestep.
- The gap between the speed of computation and speed of I/O is increasing with next generation systems.
- Developing methods to enable efficient, in situ and in transit visualization and analysis.
- Enabling communication free and re-wind capabilities.
- Supported under the Aurora Early Science Program for Data and Learning.

Planned first Exascale System in the US

Use Case 2: Ventilator Splitting

Case study: ventilator shortage

- Ventilators are vital equipment for assistance with respiration
- A shortfall of 45,000-160,000 ventilators with the U.S. was predicted for the ongoing pandemic
- Duke researchers paired up with restor3D to develop a safe, efficacious ventilator splitter and resistor system (VSRS)

Note: Ventilator splitting should **NOT** be standard-of-care and should only be used as a last recourse.

What was needed?

- Tuning to select optimal restrictor for specific patients
- Computational fluid dynamics model to capture airflow
- Validation of fluid model
- Quick time-to-solution

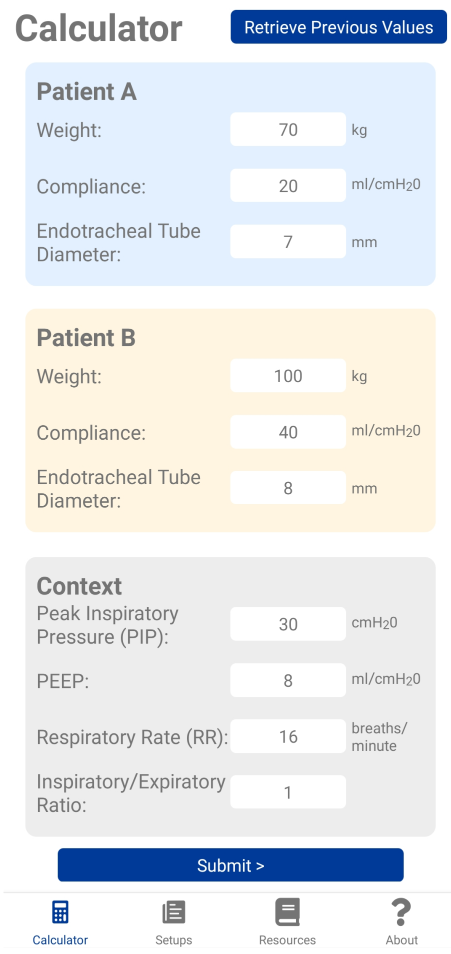

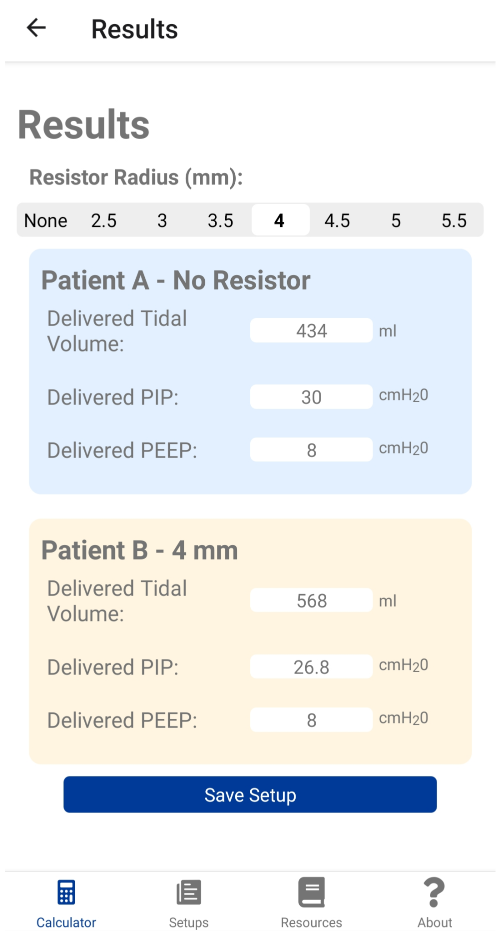

Courtesy of CrossComm

Validating Computational Model

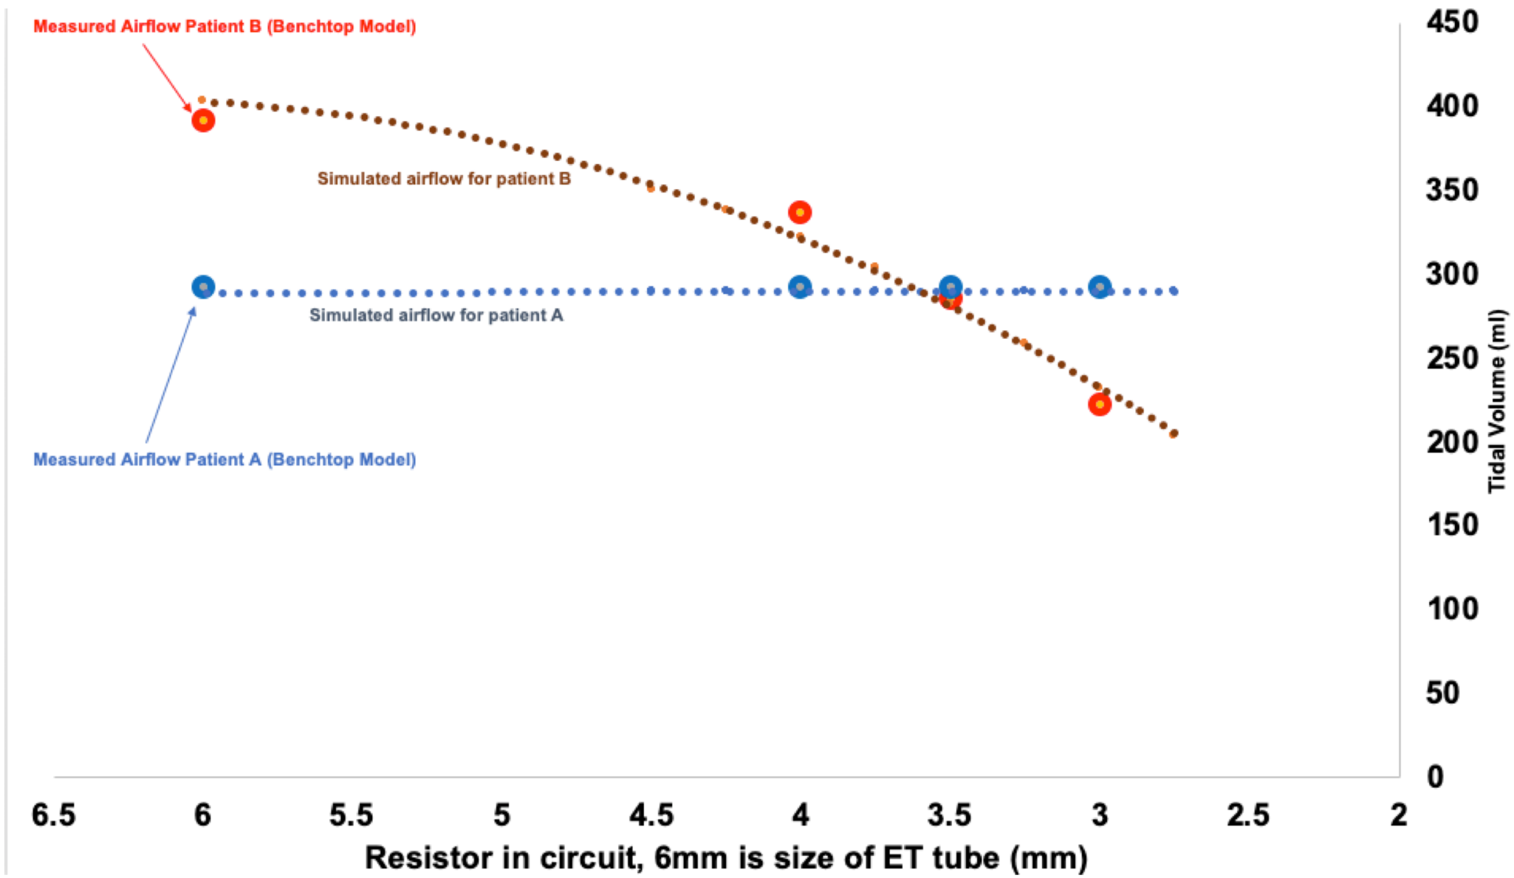

>200,000 Simulations Needed

	Minimum Value	Maximum Value	Step Size
Pulmonary Compliance (ml / cmH_2O)	10	100	1-2
Endotracheal Tube Diameter (mm)	6	8.5	0.5
Peak Inspiratory Pressure (cmH_2O)	20	50	1
Positive End-Expiratory Pressure (cmH_2O)	5	20	1
Inspiratory to Expiratory Ratio	1:3	1:1	fractional
Respiratory Rate (breaths / minute)	10	30	1
Resistor Radii (mm)	2.5	5.5	0.5

Science-Driven Solution

- COVID-19 HPC Consortium
- Coupled with Microsoft and Duke Office of Information Technology
- Worked with Microsoft Azure
- Need for minimal time-to-solution

Within days of submitting our request, we completed **800,000** compute hours in one weekend.

Kaplan et al. CiSE 2020

ACKNOWLEDGEMENTS

Randles Lab:

- **Adebayo Adebiyi**
- **Greg Herschlag**
- Mike Kaplan
- **John Gounley**
- Mahsa Dabagh
- Luiz Hegele
- **Jeff Ames**
- Simba Chidyagwai
- **Brad Feiger**
- Ismael Perez
- **Daniel Puleri**
- Sayan Roychowdhury
- **Maddie Vardhan**
- **Harvey Shi**
- **Peter Balogh**
- Zach Bernstein
- Dian Niu
- Erick Lorenzana-Saldivar
- Arpita Das
- Cyrus Tanade

Collaborators:

- James Chen, PhD, UC Denver
- Jane Leopold, MD, BWH
- Andy Kahn, MD, PhD, UCSD
- Franziska Michor, PhD, Dana Farber Cancer Institute
- Erik Draeger, PhD, LLNL
- Seyong Lee, PhD, ORNL
- Jeff Vetter, PhD, ORNL
- Joe Insley, PhD, ANL
- Silvio Rizzo, PhD, ANL
- David Frakes, PhD, ASU
- Roger Kamm, PhD, MIT
- Manesh Patel, MD, Duke
- Piers Barker, MD, Duke
- Leila Mureebe, MD, Duke
- Sreek Vemulapalli, MD, Duke
- Lynn Koweek, MD, Duke
- Schuyler Jones, MD, Duke
- Kevin Hill, MD, Duke
- Chad Hughes, MD, Duke

NIH DIRECTOR'S
EARLY
INDEPENDENCE
AWARD

Research supported by the Office of the Director, National Institutes of Health under Award Number DP5OD019876, the Duke Quantitative Initiative, the Hartwell Foundation, Boston Scientific, and the LLNL Laboratory Directed Research and Development (LDRD) program. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Computing support for this work came from the Lawrence Livermore National Laboratory (LLNL) Institutional Computing Grand Challenges.