Healthcare with Real and Virtual Sensors using AI

Prof. Jorge Ortiz

Rutgers University

Cyber-Physical Intelligence / WINLAB

Smart Healthcare

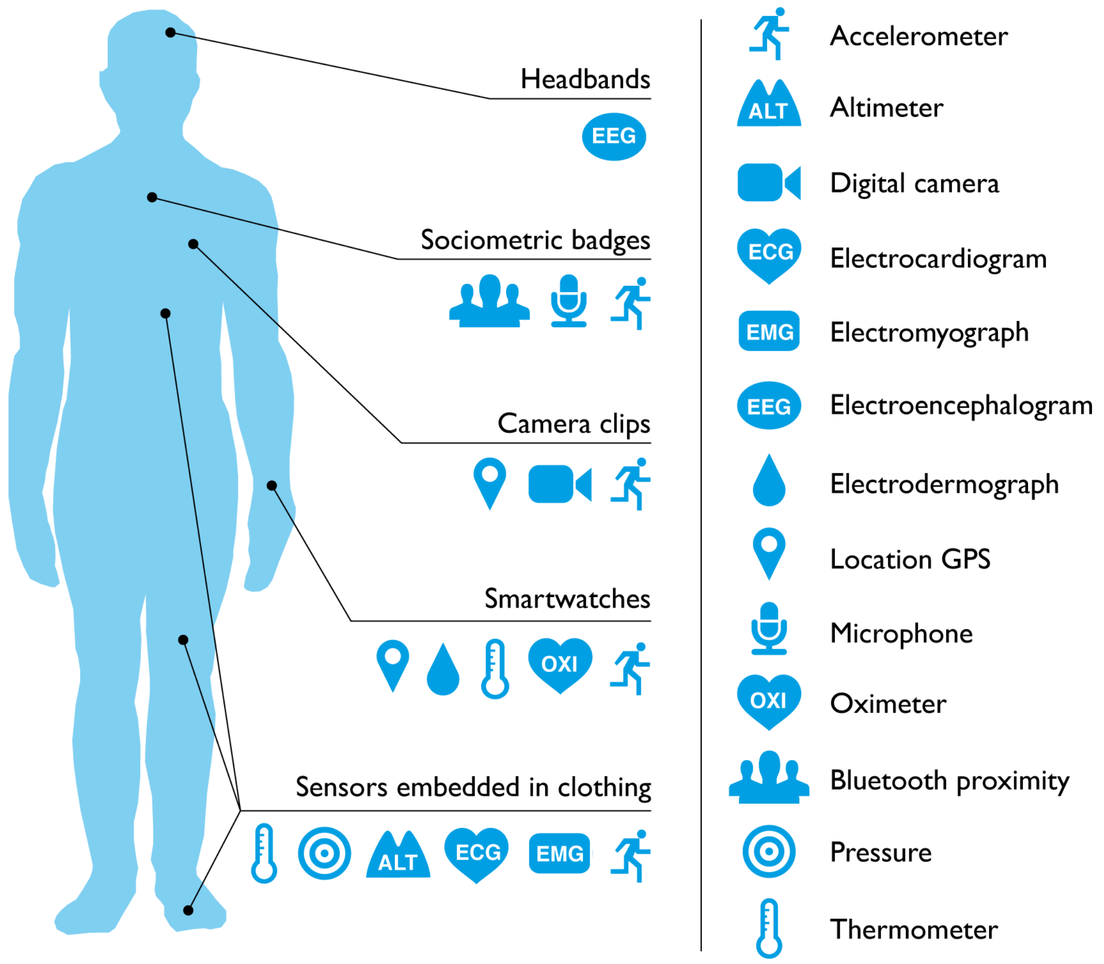

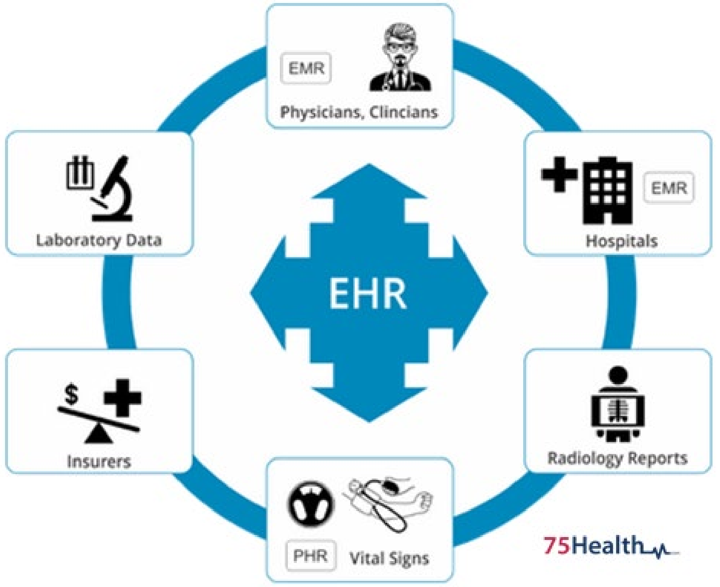

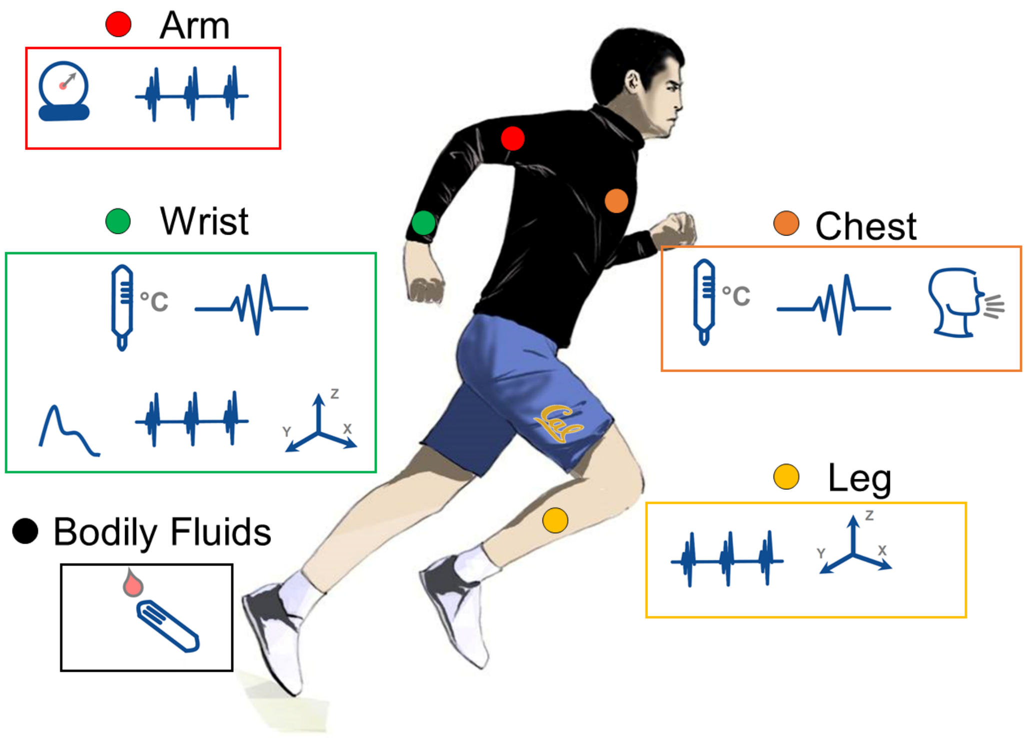

Motivation

- Between 2006 and 2030, the U.S. population of adults **aged 65+ will nearly double** from 37 million to 71.5 million people *
- **87% of adults** age 65+ want to stay in their current home and community as they age *

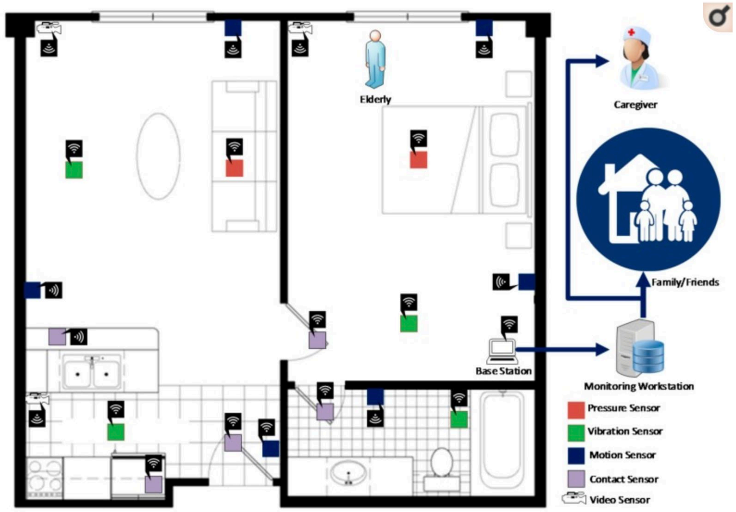

[* https://www.aarp.org/livable-communities/info-2014/livable-communities-facts-and-figures.html](https://www.aarp.org/livable-communities/info-2014/livable-communities-facts-and-figures.html)

mHealth and Connected Health: People, Technology, Process

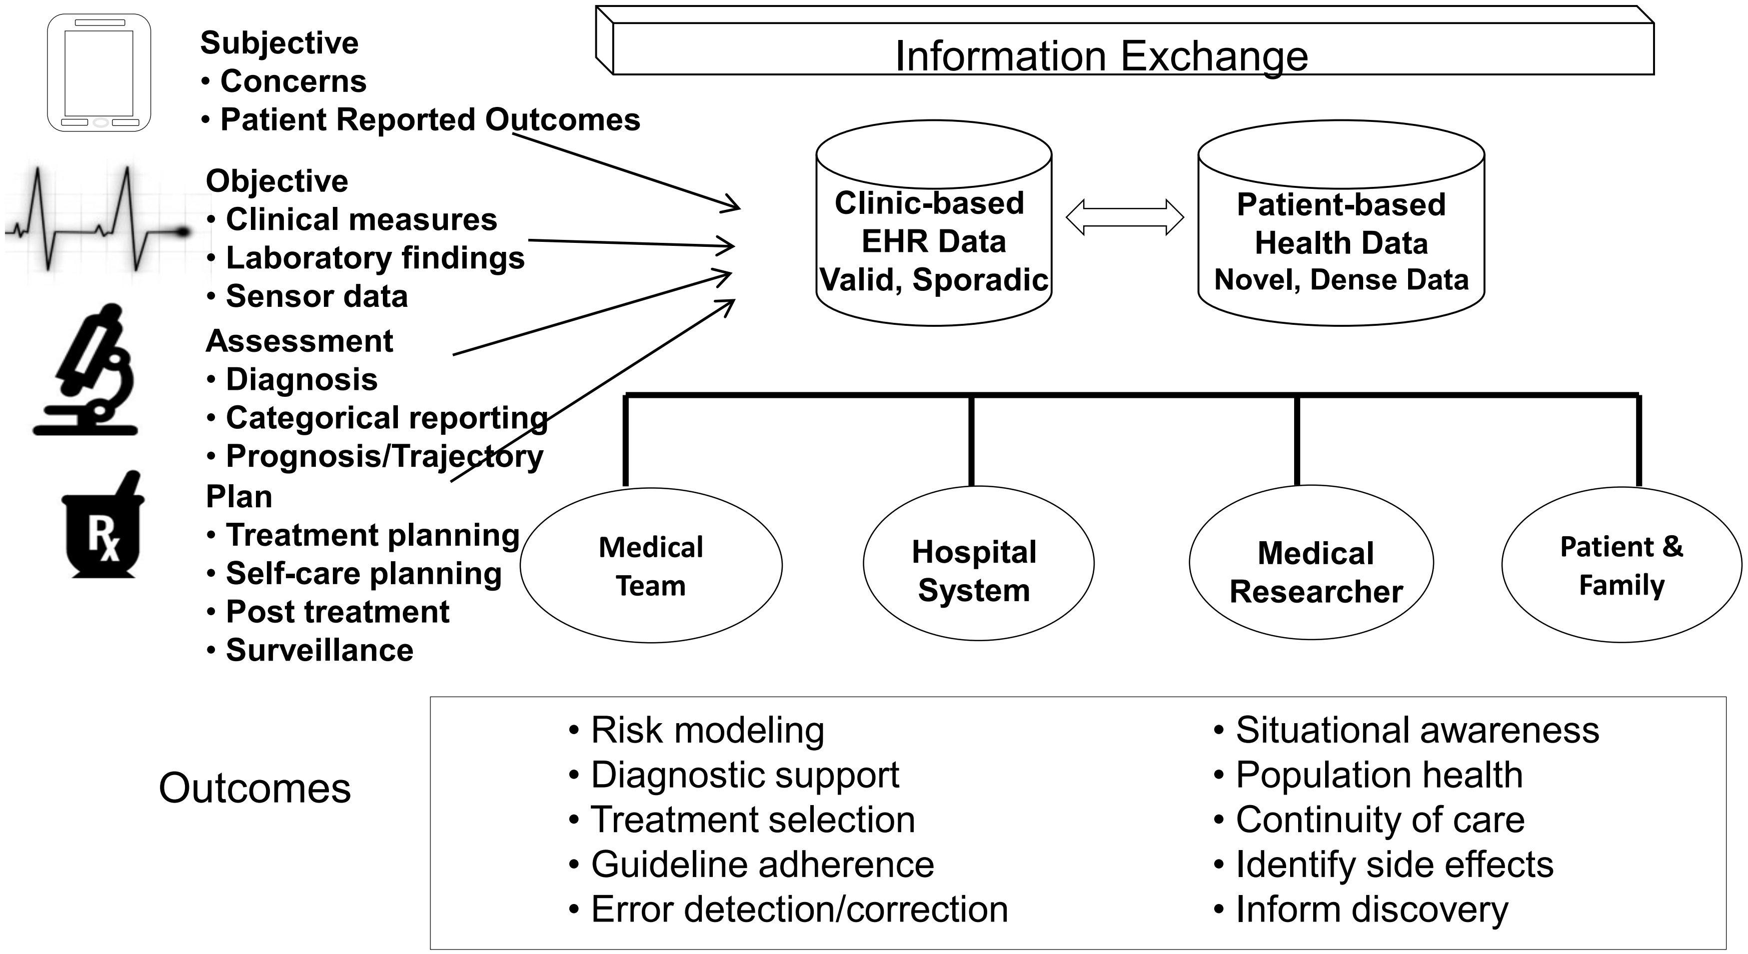

Real Sensors

Smart Healthcare

Activity monitoring for disease progression monitoring and safety

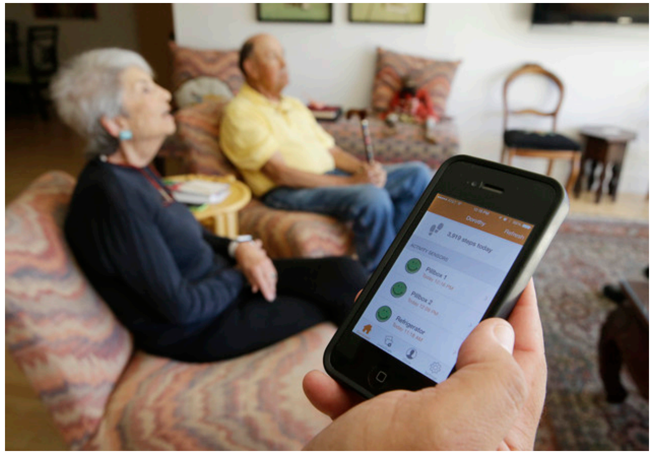

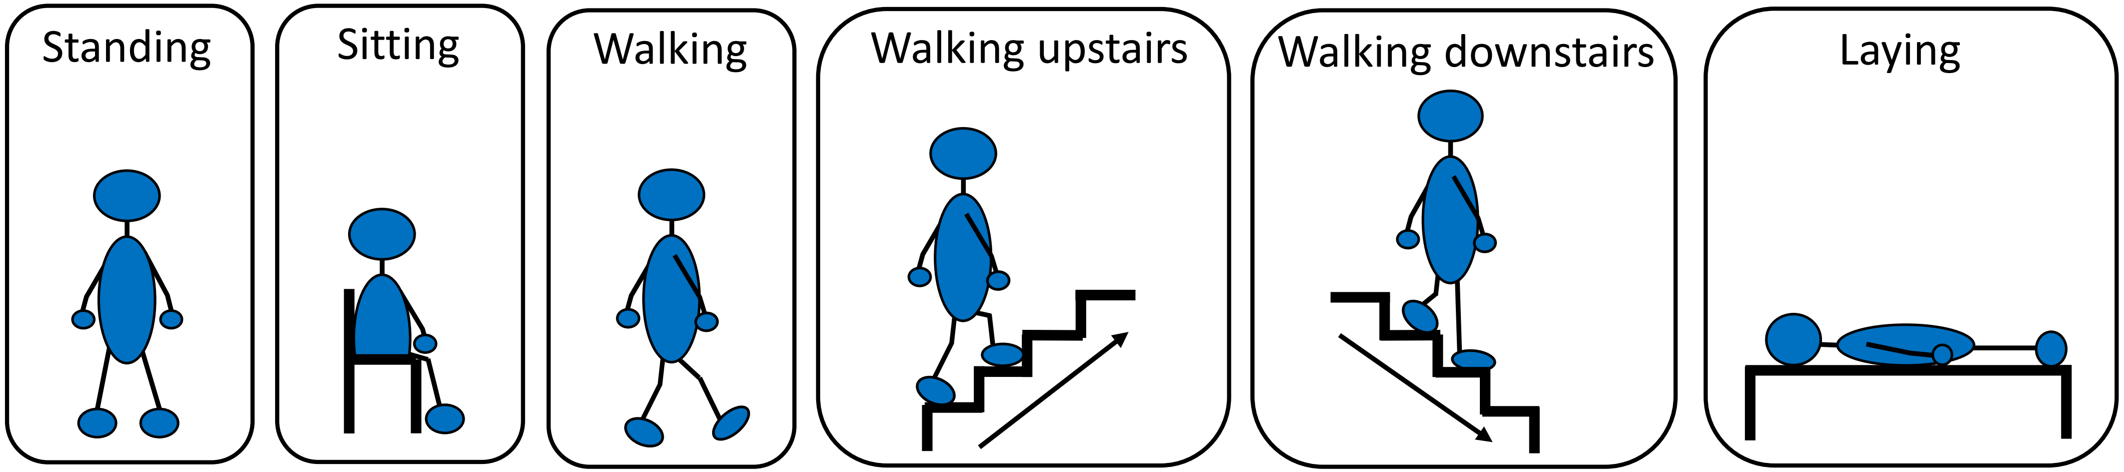

Human Activity Recognition

- Use a set of sensors and/or cameras to classify movements as they occur
- Entertainment/Gaming
- Security/Safety
- Healthcare
  - Gait analysis
  - Remote monitoring in eldercare

Labeled points of interest from RGB-D camera + trajectory analysis

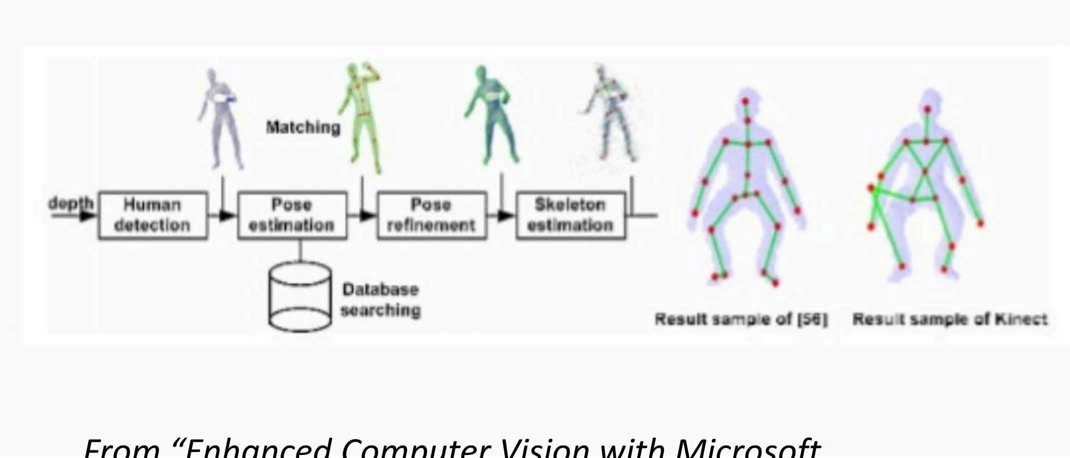

From "Enhanced Computer Vision with Microsoft Kinect Sensor: A Review"

Phillips "DirectLife" sensor

HAR w/Mobile Phones' IMU

- 561 features extracted from Mobile phone IMU stream
- Features: statistical summaries over windows readings
- **Observations**:
  - Features not independent → live on low-dimensional manifold
  - Privacy & latency:
    - Too much data to run processing on the mobile itself
    - Concern over sending data to cloud

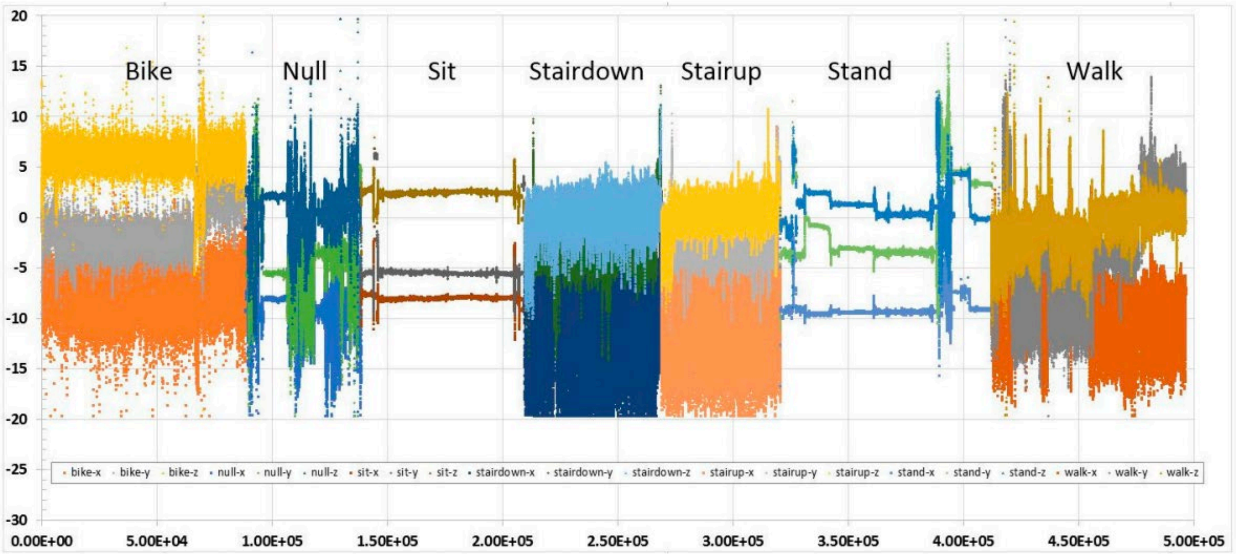

Our Approach

- Reduce to 6 key, raw IMU signals
- Image generation from multivariate time series data

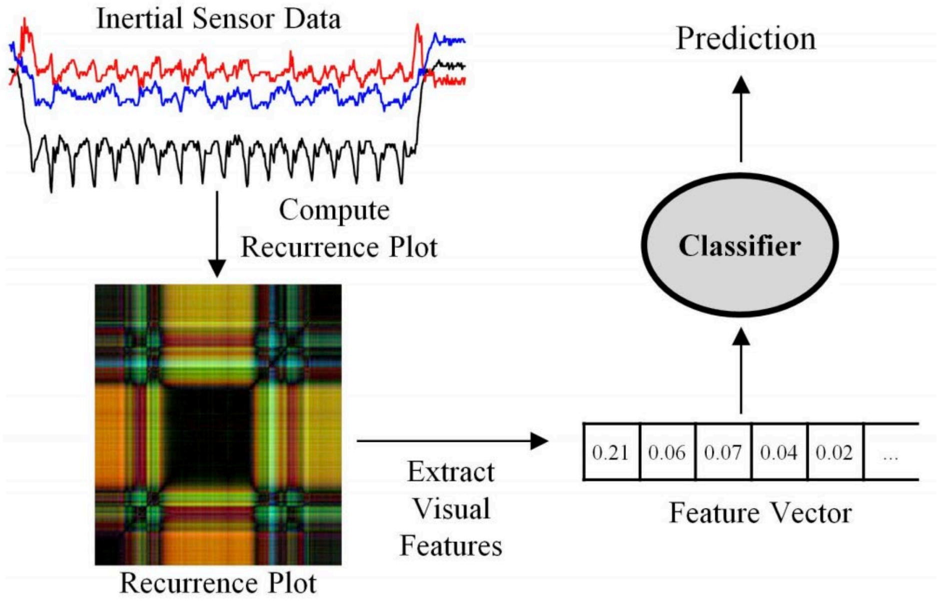

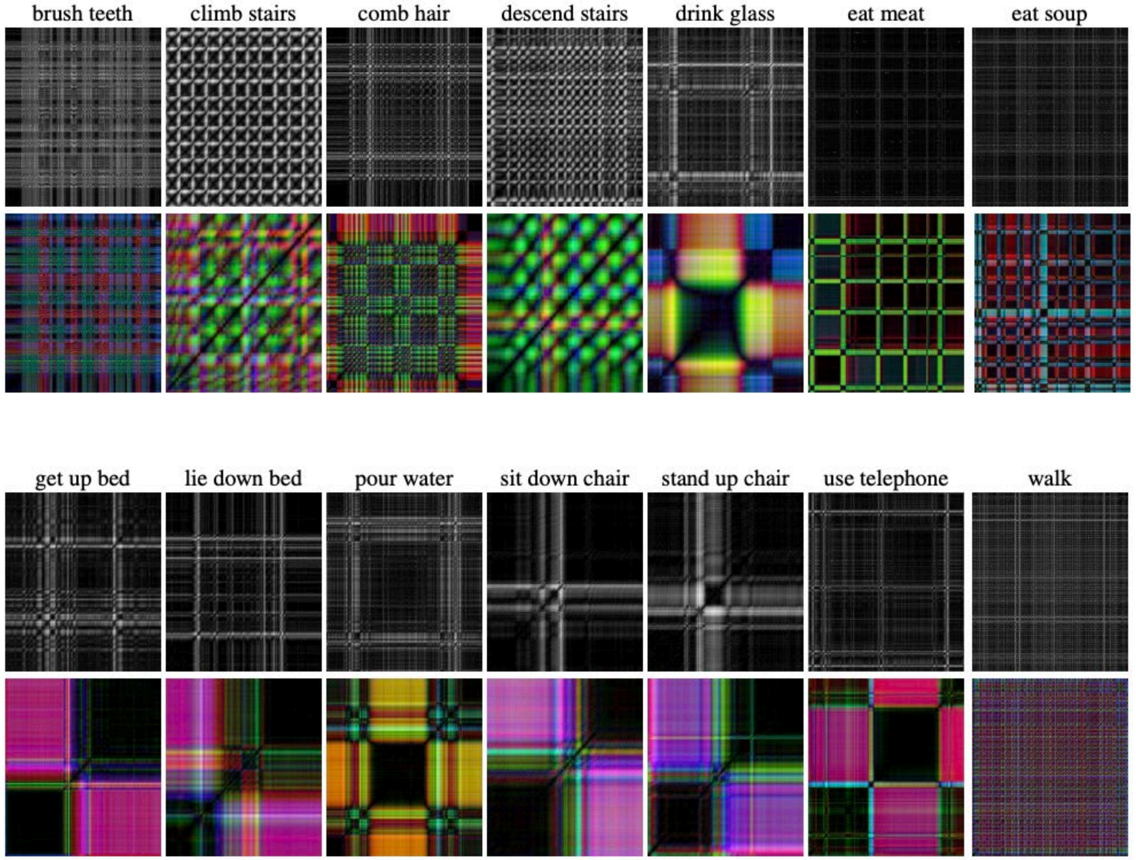

O. A. Penatti and M. F. Santos, "Human activity recognition 2018 International Joint Conference on Neural Networks (IJCNN) from mobile inertial sensors using recurrence plots," arXiv preprint arXiv:1712.01429, 2017.

Our approach: Image Classification

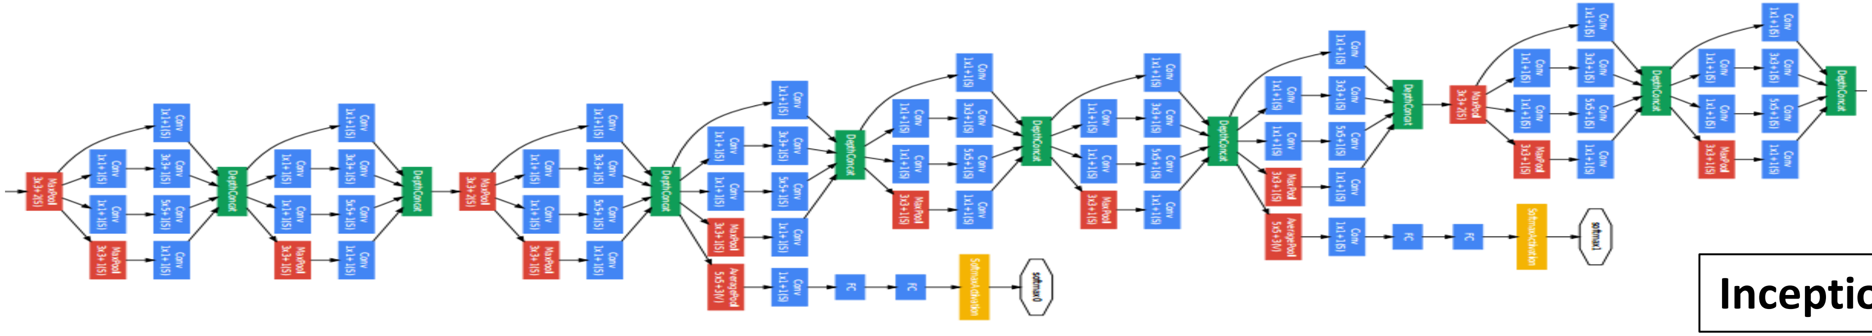

Inception

Objectives: Small and Accurate

We only use 6 accel/gyro signals since Linear accel is just total accel – gravity … e.g. redundant from an information standpoint

We only use raw signals for our image, since we found frequency space to not affect accuracy

Since our input signal image has fewer rows, our DCNN can be relatively shallow, one convolutional and one subsampling layer

We remove the Fully-Connected layer! It reduces the size of the network by 95% and also eliminates the large MxM. We found empirically it does not affect accuracy. (we use dropout during training to prevent overfitting)

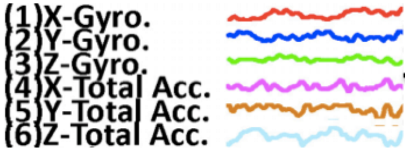

6 IMU Signals

Smart Healthcare

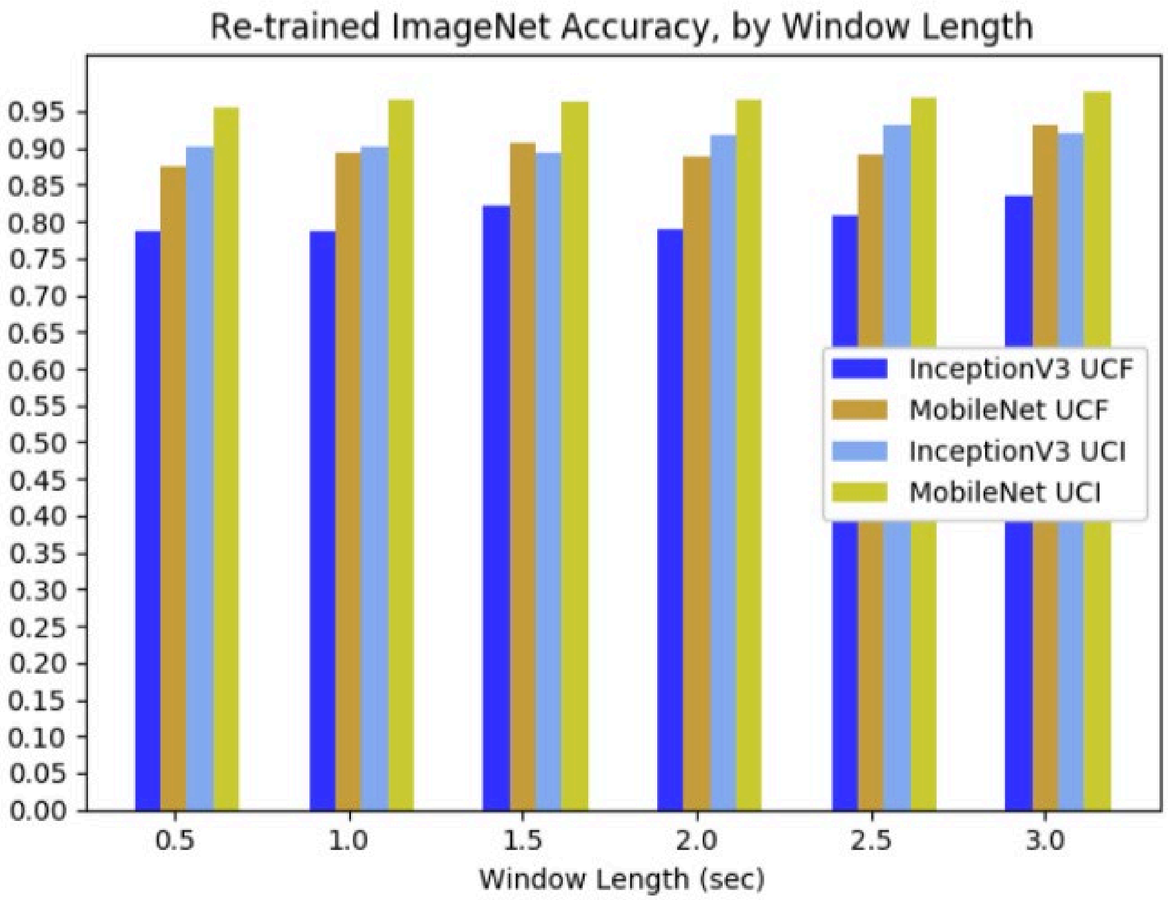

Re-trained from Scratch

Accuracy (%)	Classifier
99.93	Our DCNN+SVM HAR pipeline on 6 IMU signals
99.5	Our DCNN using 9 IMU signals
97.6	Deep CNN + SVM
96.0	Multiclass SVM
95.1	Deep CNN
93.4	Retrained Inception
91.4	LSTM-HAR

Using Transfer Learning

Virtual Sensors

CENTRAL ILLUSTRATION: Association Between Increasing Levels of Leisure-Time Physical Activity and Risk of Different Heart Failure Phenotypes

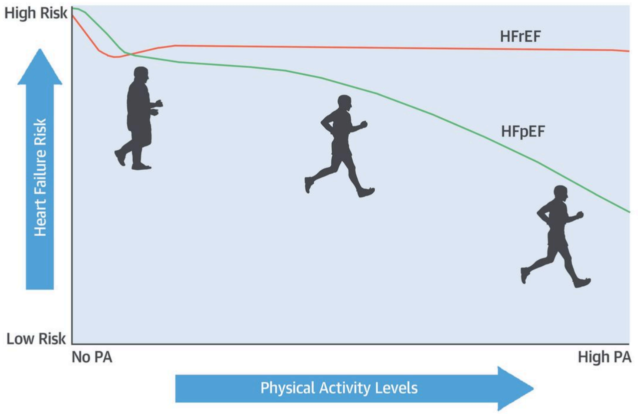

	Guideline Recommended Minimum PA	2 x Guideline Recommended Minimum PA	3 x Guideline Recommended Minimum PA
Brisk Walking ~3.35 METs	150 minutes per week	300 minutes per week	450 minutes per week
Jogging/Running ~6.5-7 METs	75 minutes per week	150 minutes per week	225 minutes per week

Pandey, A. et al. J Am Coll Cardiol. 2017;69(9):1129-42.

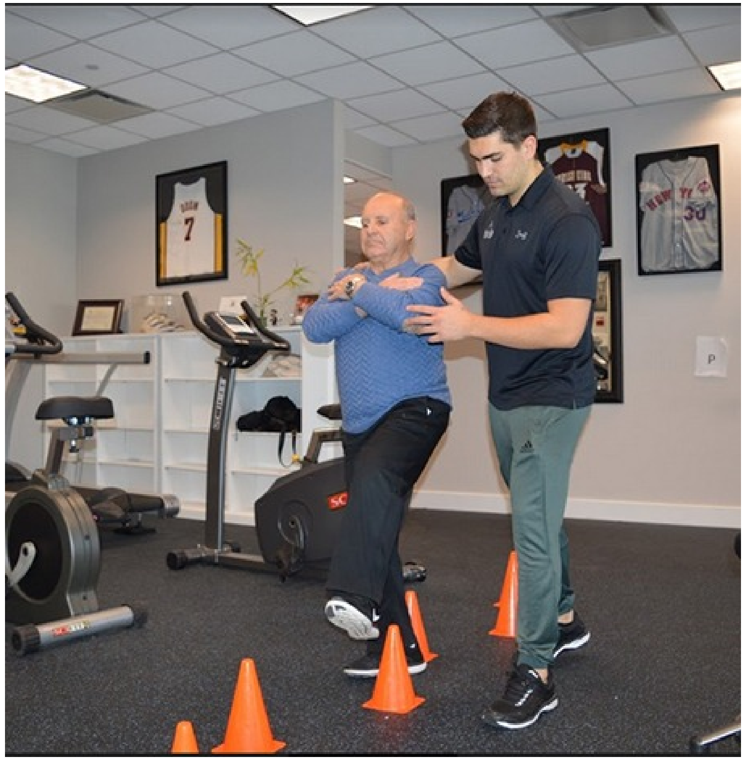

Generative Models for IMU Data from Video

- Real sensors have many limitations for monitoring and significantly reduce **quality of life in the very elderly**
- No requirement that sensing system has to be comfortable to be approved
- Track body movements from video and **generate synthetic IMU sensor** data from it
- Uses **deep-learning based keypoint tracking** from the video.

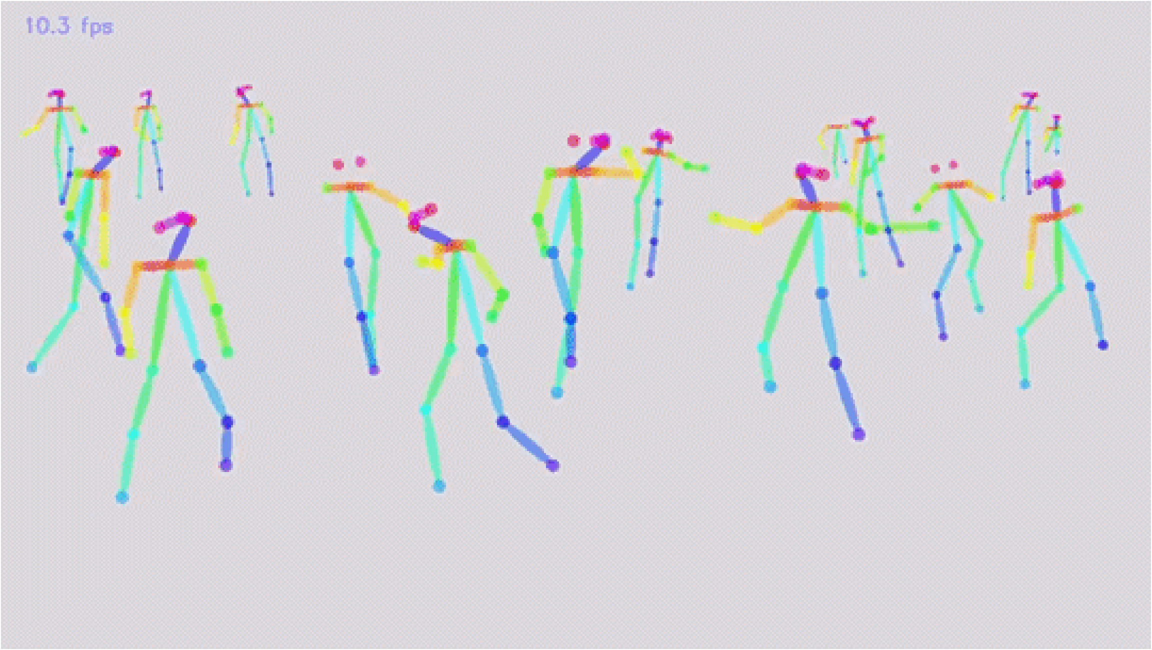

* Cao and Tomas Simon and Shih-En Wei and Yaser Sheikh, Realtime Multi-Person 2D Pose Estimation using Part Affinity FieldsZhe, CVPR 2017

Progress

- Single person motion track

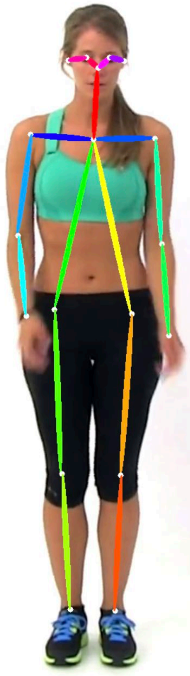

joint	nose	neck	right-shoulder	right-elbow	right-wrist
person 0	(944.0,185.0)	(937.0,287.0)	(847.0,280.0)	(829.0,425.0)	(841.0,540.0)

left-shoulder	left-elbow	left-wrist	right-hip	right-knee	right-ankle
(1025.0,286.0)	(1035.0,437.0)	(1045.0,568.0)	(881.0,531.0)	(892.0,766.0)	(904.0,960.0)

left-hip	left-knee	left-ankle	right-eye	left-eye	right-ear	left-ear
(992.0,537.0)	(976.0,768.0)	(968.0,962.0)	(927.0,170.0)	(960.0,170.0)	(901.0,174.0)	(981.0,175.0)

Track one person with eighteen joints movements in a video.

Progress

- Single person motion track

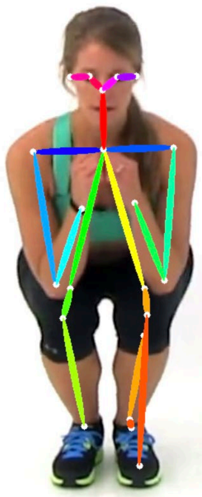

joint	nose	neck	right-shoulder	right-elbow	right-wrist
person 0	(929.0,560.0)	(930.0,632.0)	(845.0,635.0)	(873.0,796.0)	(904.0,704.0)
person 1	(None, None)	(None, None)	(None, None)	(None, None)	(None, None)

left-shoulder	left-elbow	left-wrist	right-hip	right-knee	right-ankle
(1015.0,630.0)	(1004.0,791.0)	(968.0,697.0)	(890.0,801.0)	(881.0,837.0)	(907.0,968.0)
(None, None)	(None, None)	(None, None)	(None, None)	(None, None)	(None, None)

left-hip	left-knee	left-ankle	right-eye	left-eye	right-ear	left-ear
(980.0,801.0)	(983.0,833.0)	(974.0,1016.0)	(913.0,543.0)	(947.0,544.0)	(890.0,544.0)	(970.0,543.0)
(979.0,858.0)	(962.0,960.0)	(964.0,970.0)	(None, None)	(None, None)	(None, None)	(None, None)

Track one person with eighteen joints movements in a video.

Current Status

- Test and calculate specific joint movement in a single person video

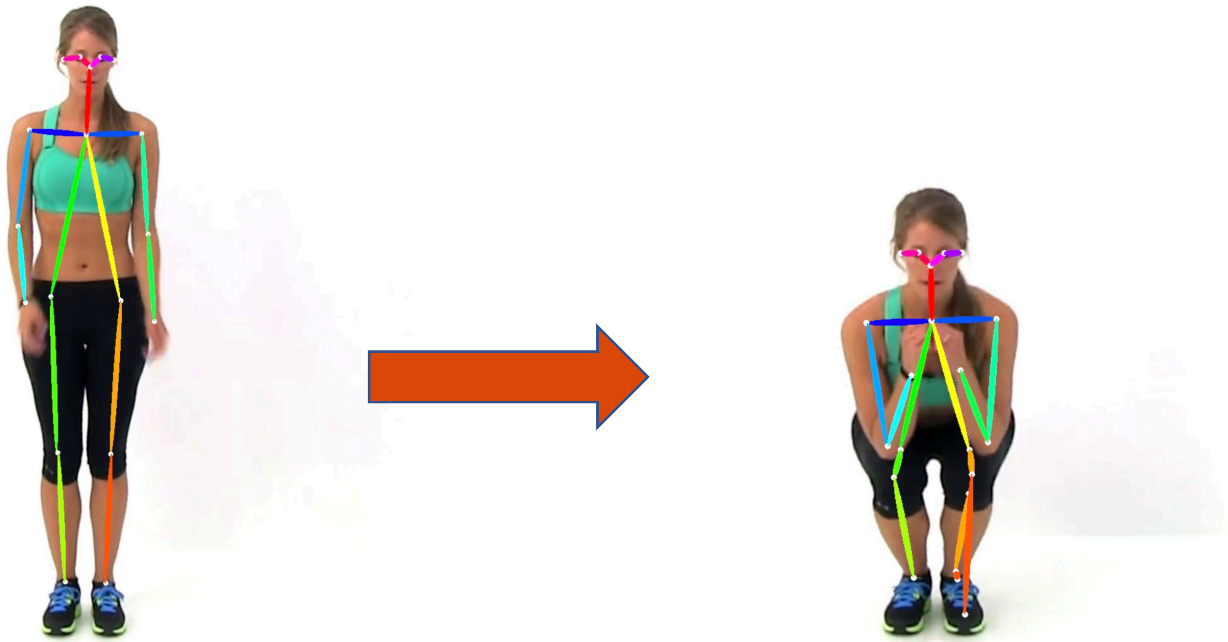

```
The left shoulder joint starts moving in ([1025.0, 286.0])
The left shoulder joint ends moving in ([1015.0, 630.0])
Horizontal direction: -10.0
Vertical direction: 344.0
```

Track and calculate left shoulder joint position movement in the squat action as an example.

Progress

- Multi-person motion track

joint	nose	neck	right-shoulder
person 0	(882.0,408.0)	(863.0,458.0)	(817.0,459.0)
person 1	(1476.0,408.0)	(1473.0,442.0)	(1441.0,442.0)
person 2	(1796.0,382.0)	(1790.0,438.0)	(1740.0,440.0)
person 3	(1341.0,409.0)	(1333.0,473.0)	(1290.0,467.0)
person 4	(1030.0,440.0)	(1032.0,465.0)	(1008.0,468.0)
person 5	(1163.0,425.0)	(1160.0,457.0)	(1135.0,456.0)
person 6	(691.0,425.0)	(692.0,460.0)	(660.0,454.0)
person 7	(1555.0,408.0)	(1556.0,448.0)	(1521.0,449.0)
person 8	(524.0,423.0)	(528.0,470.0)	(451.0,476.0)
person 9	(778.0,428.0)	(775.0,458.0)	(748.0,459.0)
person 10	(132.0,379.0)	(135.0,442.0)	(76.0,443.0)
person 11	(390.0,444.0)	(318.0,470.0)	(363.0,482.0)
person 12	(470.0,436.0)	(471.0,472.0)	(None, None)

Track the multi-person each joints movements in a video.
Each person pose composed of eighteen joints.

Imaging Geometry

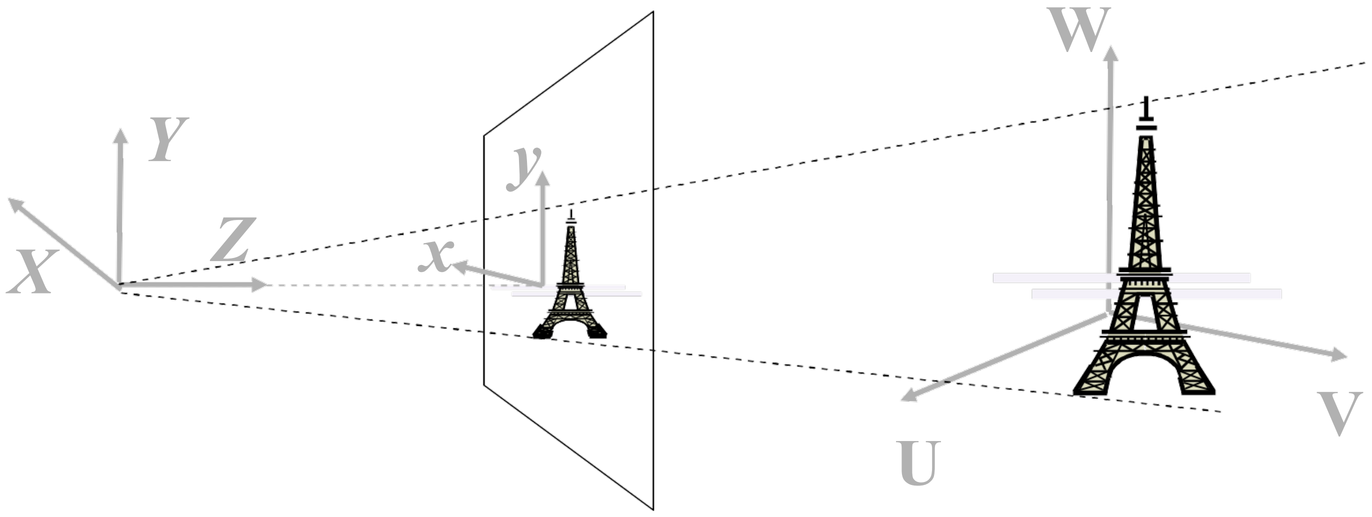

Forward Projection onto image plane.
3D (X,Y,Z) projected to 2D *(x,y)*

Imaging Geometry

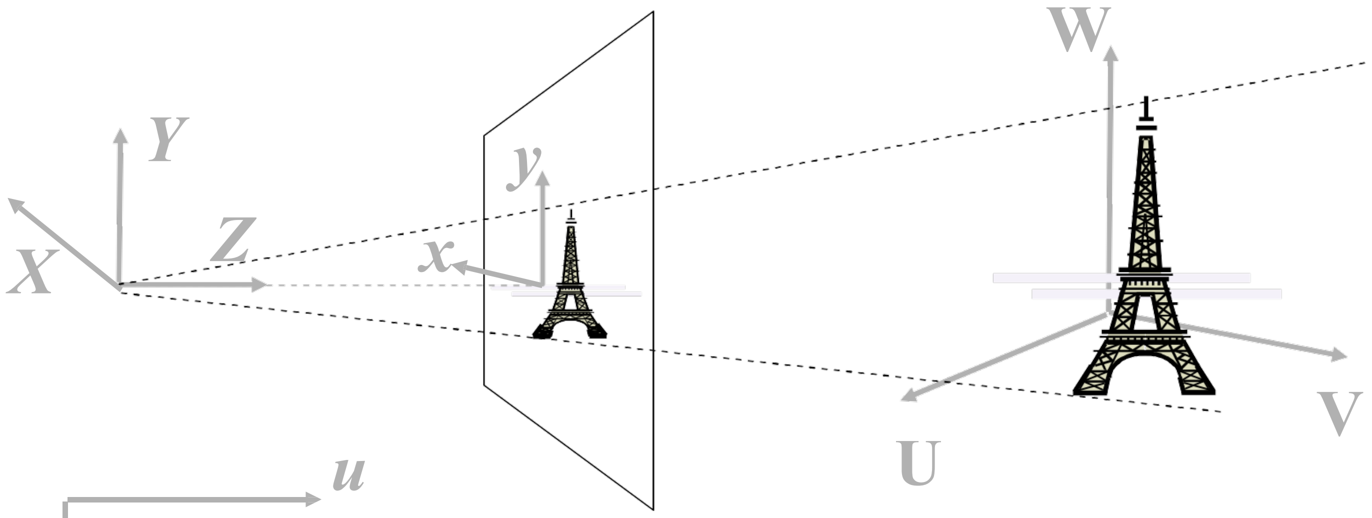

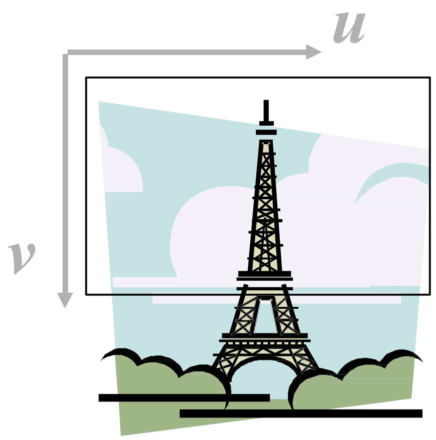

Our image gets digitized into pixel coordinates *(u,v)*

Imaging Geometry

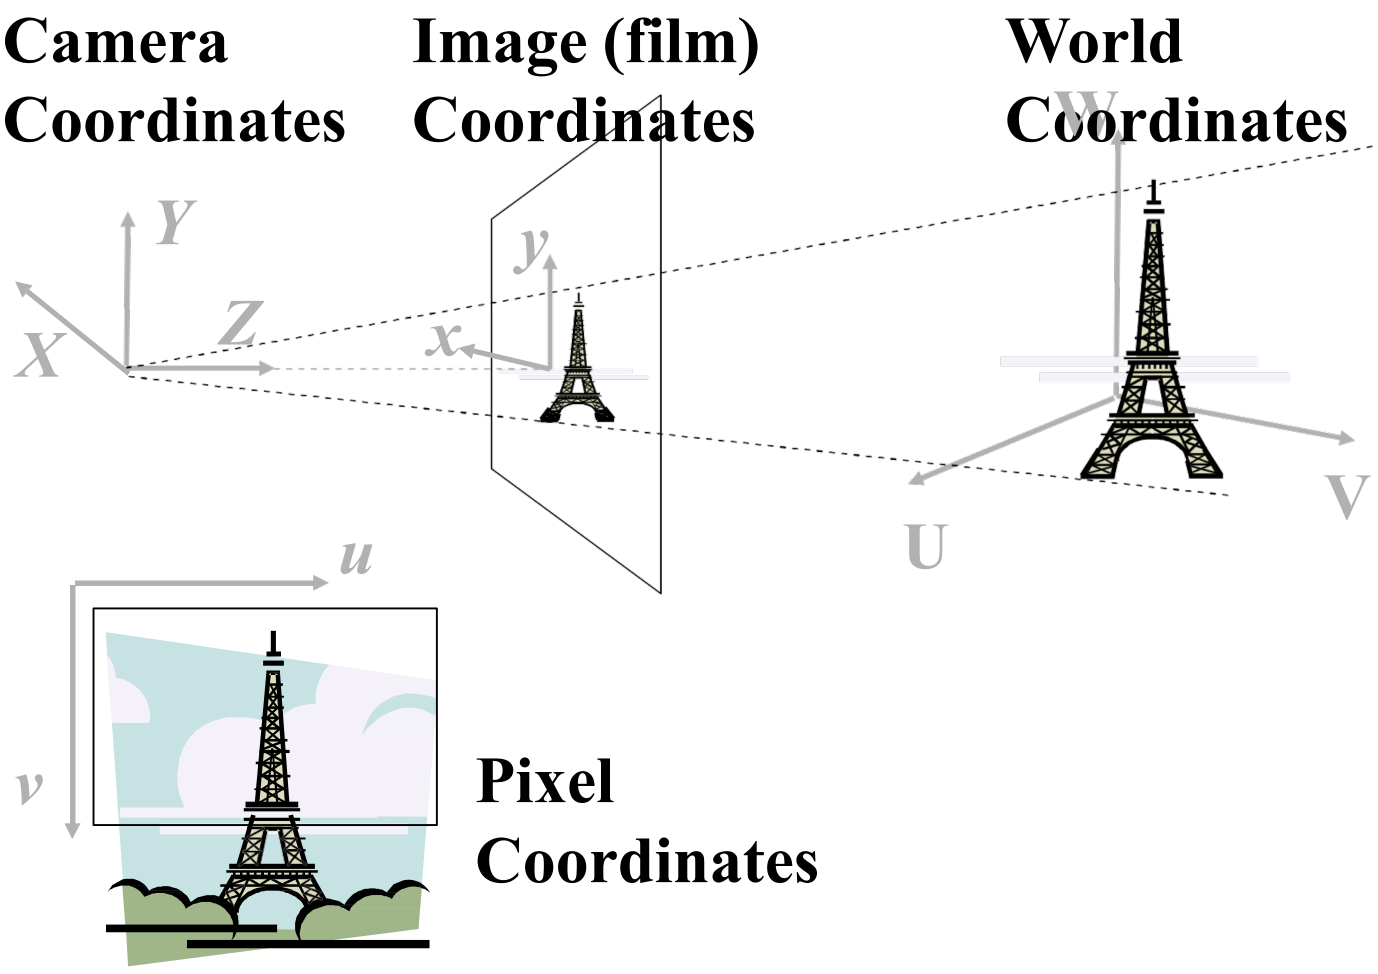

Forward Projection

World Coords		Camera Coords		Film Coords		Pixel Coords
U V W	$\Rightarrow$	X Y Z	$\Rightarrow$	x y	$\Rightarrow$	u v

We want a mathematical model to describe how 3D World points get projected into 2D Pixel coordinates.

Backward Projection

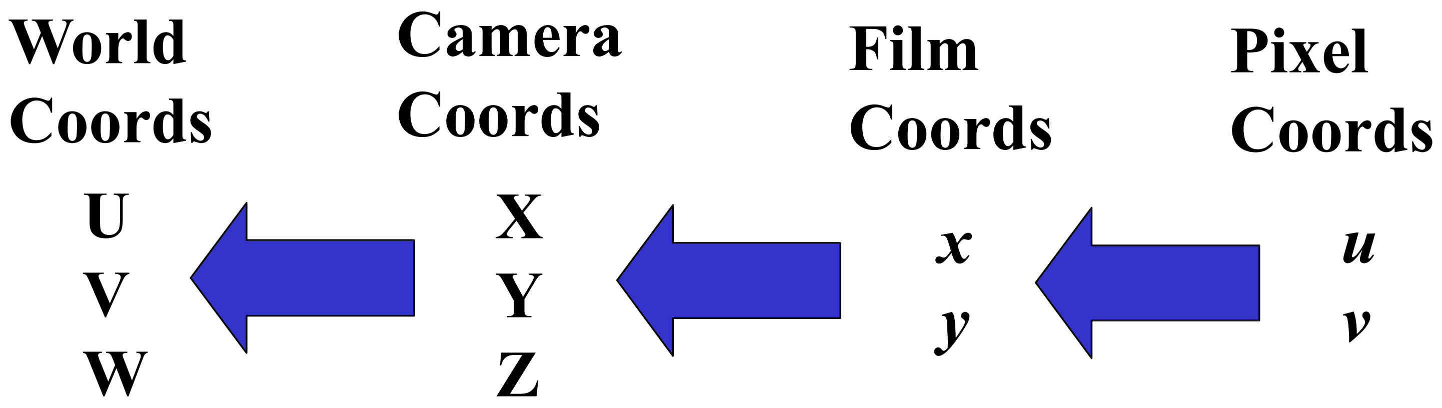

Note, much of vision concerns trying to derive backward projection equations to recover 3D scene structure from images (via stereo or motion)

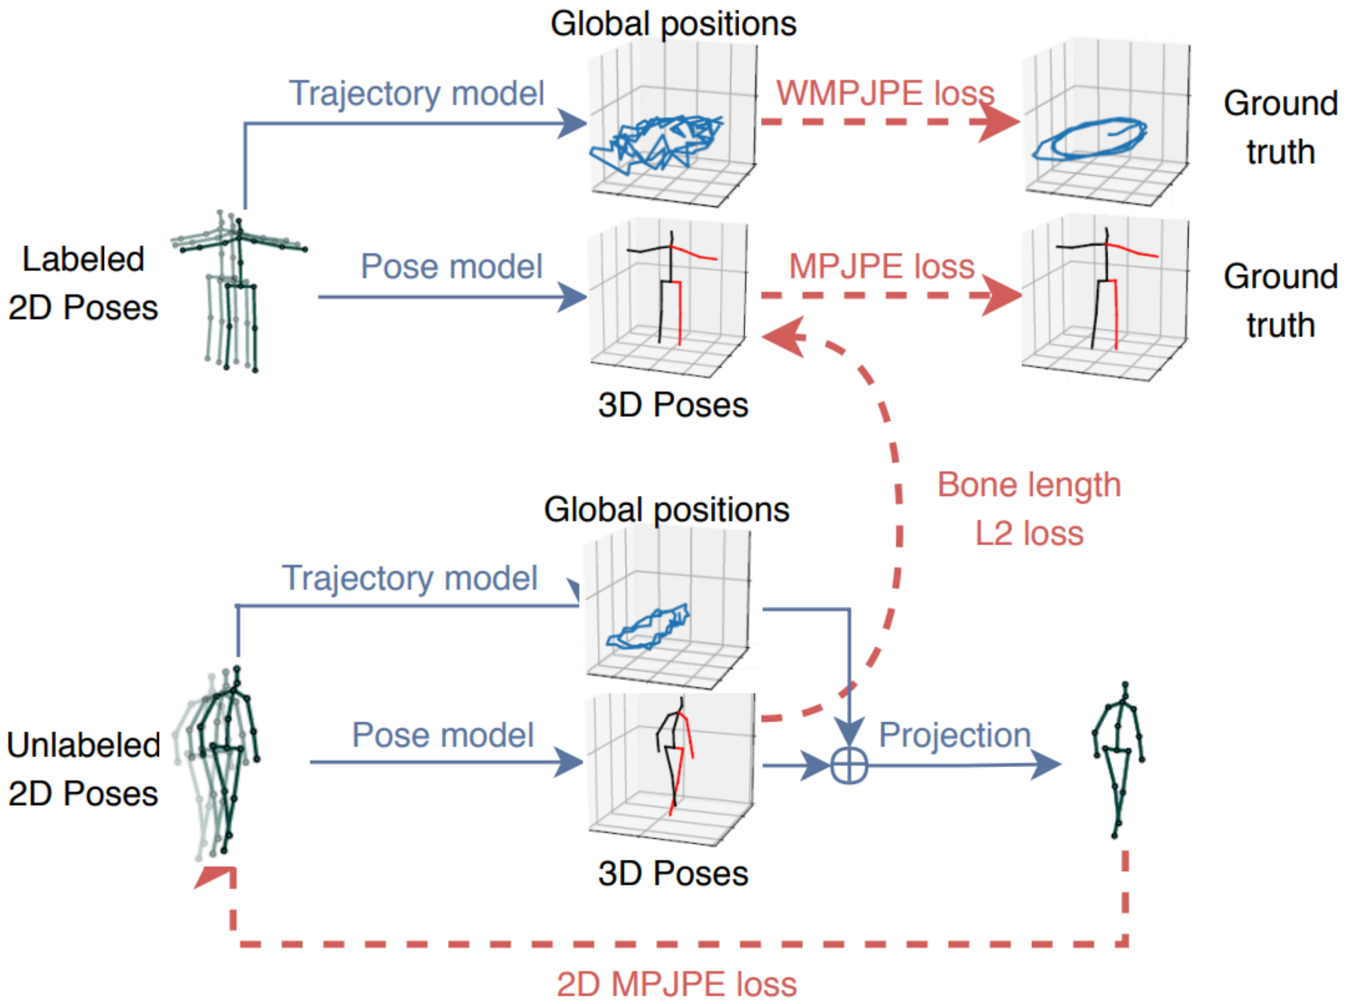
Global positions
Trajectory model
WMPJPE loss
Ground truth
Labeled 2D Poses
Pose model
MPJPE loss
Ground truth
3D Poses
Bone length L2 loss
Global positions
Trajectory model
Unlabeled 2D Poses
Pose model
Projection
3D Poses
2D MPJPE loss

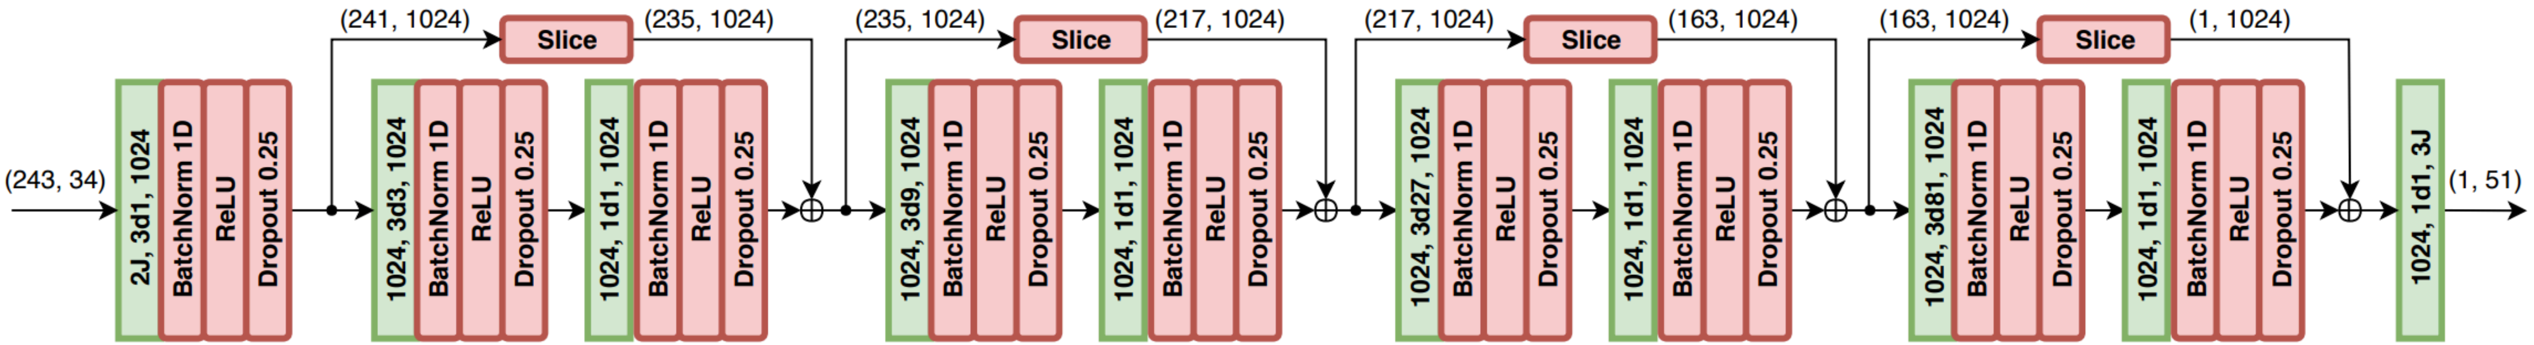
(243, 34)
2J, 3d1, 1024
BatchNorm 1D
ReLU
Dropout 0.25
(241, 1024)
Slice
(235, 1024)
1024, 3d3, 1024
BatchNorm 1D
ReLU
Dropout 0.25
1024, 1d1, 1024
BatchNorm 1D
ReLU
Dropout 0.25
(235, 1024)
Slice
(217, 1024)
1024, 3d9, 1024
BatchNorm 1D
ReLU
Dropout 0.25
1024, 1d1, 1024
BatchNorm 1D
ReLU
Dropout 0.25
(217, 1024)
Slice
(163, 1024)
1024, 3d27, 1024
BatchNorm 1D
ReLU
Dropout 0.25
1024, 1d1, 1024
BatchNorm 1D
ReLU
Dropout 0.25
(163, 1024)
Slice
(1, 1024)
1024, 3d81, 1024
BatchNorm 1D
ReLU
Dropout 0.25
1024, 1d1, 1024
BatchNorm 1D
ReLU
Dropout 0.25
1024, 1d1, 3J
(1, 51)

Motion evaluation and pose stability assessments

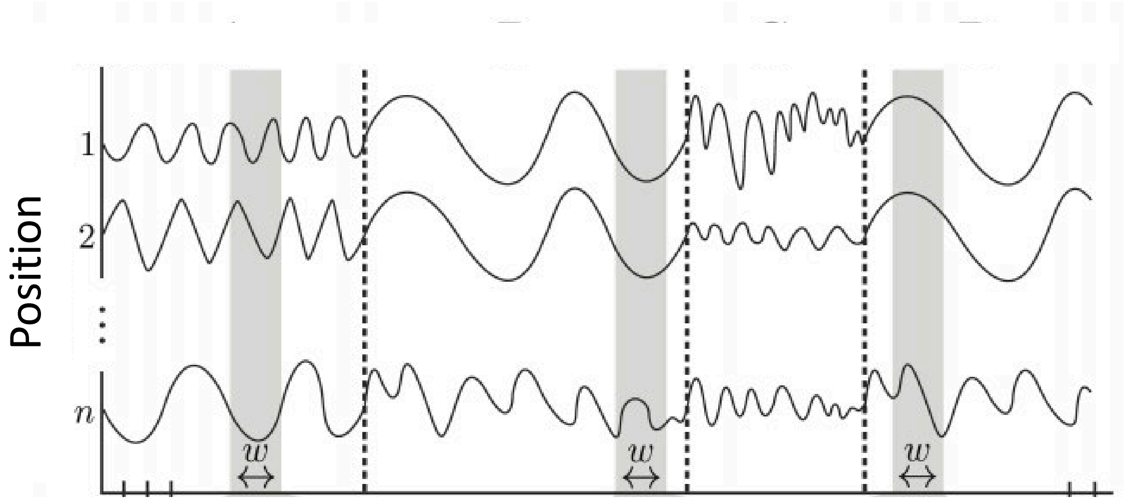

- Generate multivariate time series of positions
- Cluster them
- Look at evolution of clusters
  - Find outliers, compare similar poses, transitions, etc.
- Identify risks, compare players

Sustainability & IoT

CONTACT:

Jorge Ortiz
jorge.ortiz@rutgers.edu
http://jorgeortizphd.info

Sustainability

- "Automated Metadata Construction to support Portable Building Applications" Buildsys 2015
- "The Building Adapter: Towards Quickly Applying Building Analytics at Scale", Buildsys 2015
- "Strip, Bind, and Search: A Method for Identifying Abnormal Energy Consumption in Buildings", IPSN 2013
- "Towards Automatic Spatial Verification of Sensor Placement in Buildings", Buildsys 2013

Internet-of-Things

- *"DeviceMein:* Network Device Behavior Modeling for Identifying Unknown IoT Devices. ACM/IEEE Conference on Internet of Things Design and Implementation 2019". To appear April 2019.
- "Time Series Segmentation Through Automatic Feature Learning", arxiv 2018
- *"Deep Learning for Real-time Human Activity Recognition with Mobile Phones",* IEEE International Joint Conference on Neural Networks IJCNN 2018


Elderly
Caregiver
Family/Friends
Base Station
Monitoring Workstation
Pressure Sensor
Vibration Sensor
Motion Sensor
Contact Sensor
Video Sensor

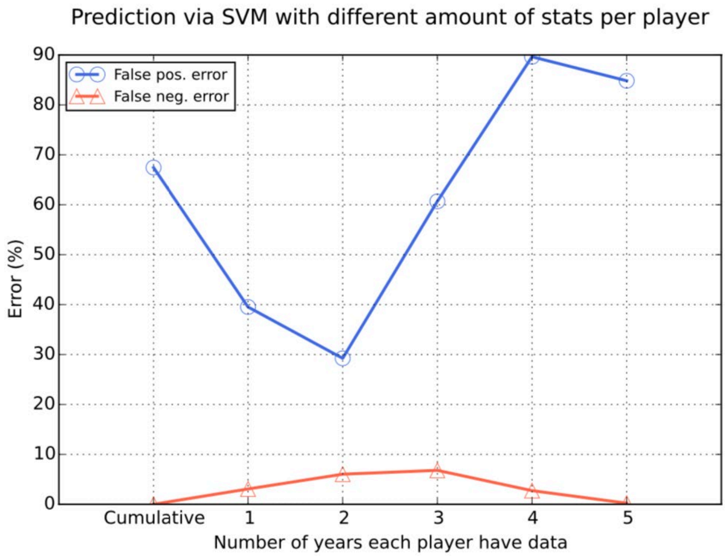

Features | **Targets**

	Game Date	Innings	ERA	Days Rest	Temperature	Injured
	4/29/16	6.0	4.00	5	42	0
	5/5/16	6.2	2.46	6	43	0
	5/11/16	7.2	6.34	6	75	0
INTERVENTION POINT →	5/18/16	2.0	12	7	55	0
	5/25/16	--	--	--	--	1

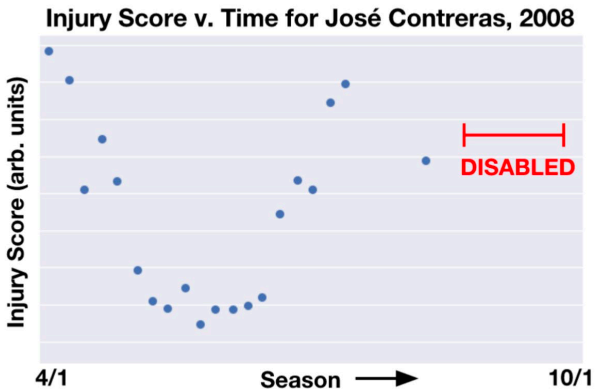

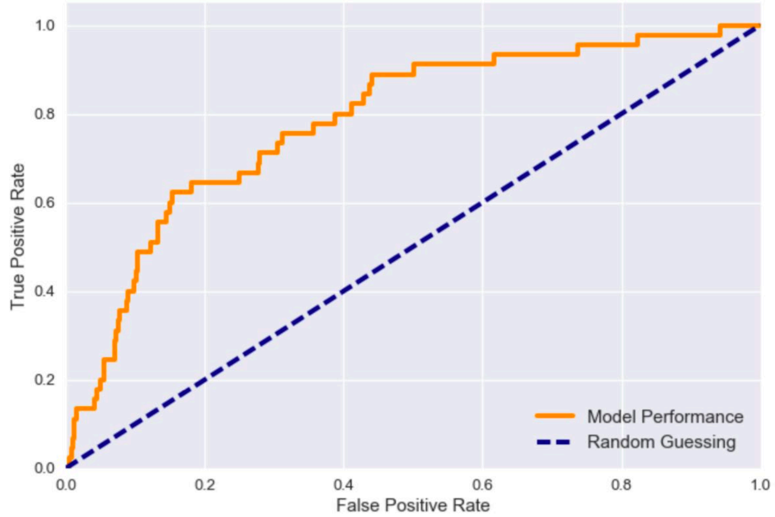

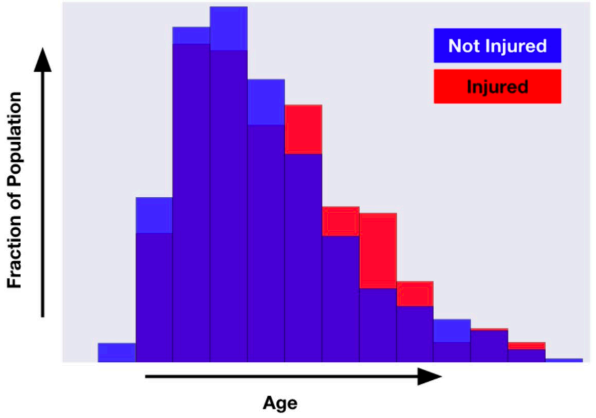

Injury Prediction: Aggregate Statistics

Injury Prediction: Biomechanics & Game Statistics

Time point	**Stride foot contact**	**Maximal shoulder external rotation**	**Ball release**

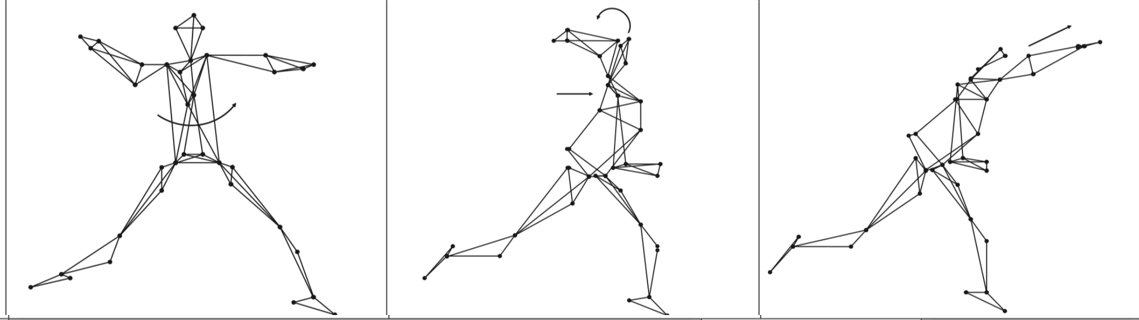

Phase	**Arm cocking**		**Acceleration**	**Deceleration**
Kinematics	Rapid upper torso rotation causes the arm to lag behind the upper torso and force the throwing shoulder into horizontal abduction	Forearm lag behind the arm and force the shoulder into external rotation (170-190°)	Rapid shoulder internal rotation (6000-7000°) and elbow extension	Deceleration of shoulder rotation

Injury	Anterior instability	Posterior impingement	SLAP lesion, posterior and subacromial impingement, growth plate injury	UCL sprain, medial epicondylitis, ulnar neuritis, stress fracture, osteochondral defect	Biceps tendonitis, rotator cuff strain, sprain, medial epicondylitis, UCL sprain	Biceps tendonitis, SLAP lesion, rotator cuff strain, subacromial impingement

[3] Baseball throwing mechanics as they relate to pathology and performance - a review.