The Kinect for Windows sensor and SDK

Dr Leandro Dihl

The Kinect for Windows

Skeletal Tracking

- Skeletal Tracking allows Kinect to recognize people and follow their actions.
- Using the infrared (IR) camera, Kinect can recognize up to six users in the field of view of the sensor. Of these, up to two users can be tracked in detail.

Skeletal Tracking

- An application can locate the joints of the tracked users in space and track their movements over time.

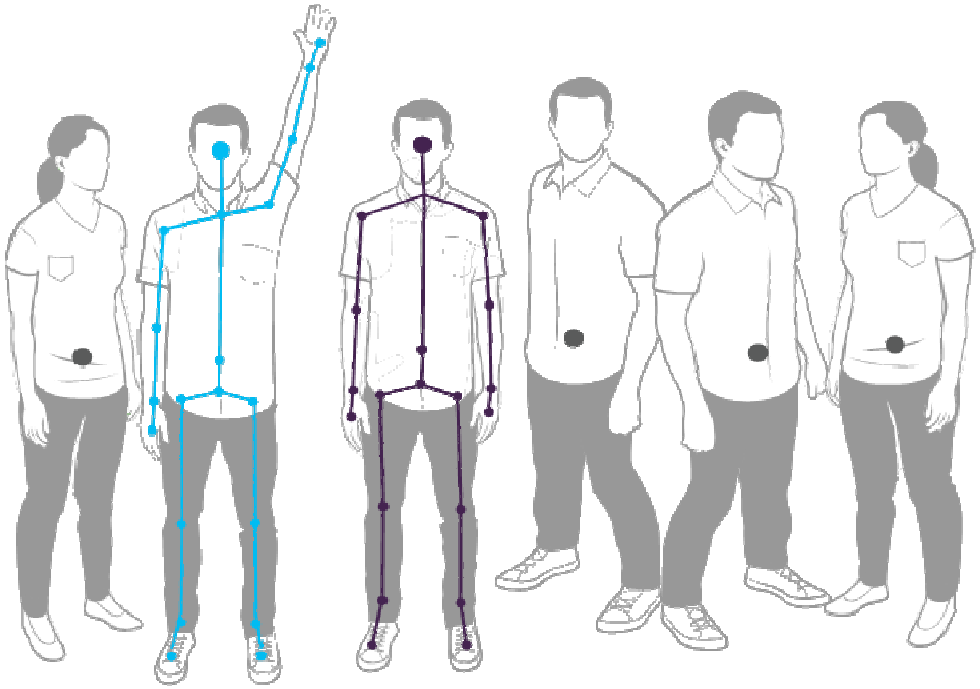

Skeletal Tracking

- Skeletal Tracking is optimized to recognize users standing or sitting, and facing the Kinect; sideways poses provide some challenges regarding the part of the user that is not visible to
- To be recognized, users simply need to be in front of the sensor, making sure the sensor can see their head and upper body; no specific pose or calibration action needs to be taken for a user to be tracked. the sensor.

Field of View

- Kinect field of view of the users is determined by the settings of the IR camera, which are set with the DepthRange Enumeration.
  - Depth data ranges, which determine how close to (or far from) the sensor a person can be and still be in the field of view.

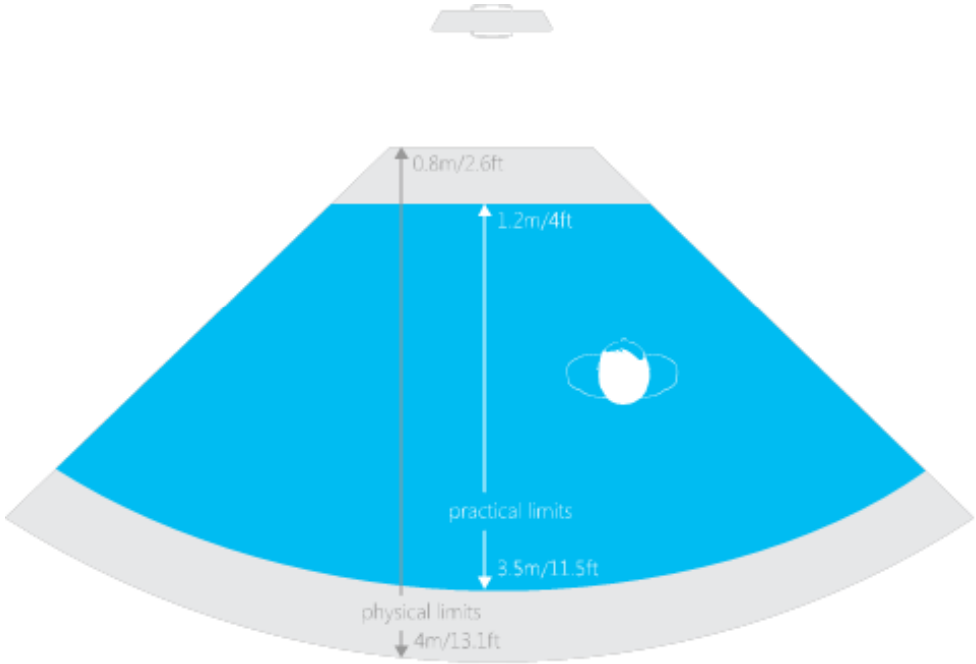

Field of View

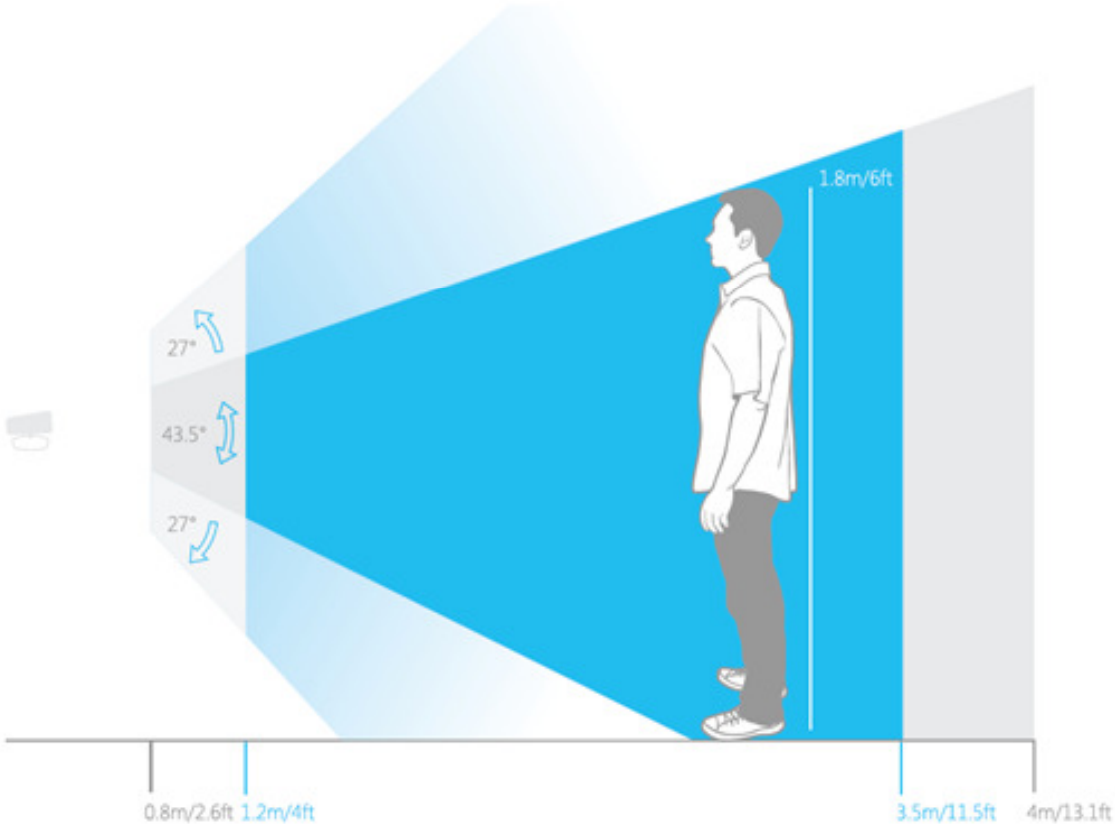

Skeletal Tracking Precision and Multiple Kinect Sensors

- The infrared emitter of a Kinect sensor projects a pattern of infrared light. This pattern of light is used to calculate the depth of the people in the field of view allowing the recognition of different people and different body parts. If you use more than one Kinect sensor to illuminate the target area, you may notice a reduction in the accuracy and precision of skeletal tracking due to interference with the infrared light sources.

Tracking Users with Kinect Skeletal Tracking

- Enabling Skeletal Tracking
  - To track users, an application needs to enable skeletal tracking.
- To enable skeletal tracking, call the INuiSensor::NuiSkeletonTrackingEnable method. To receive information about the recognized users, call the INuiSensor::NuiSkeletonGetNextFrame method.

How to access skeletal tracking information

- Every time INuiSensor::NuiSkeletonGetNextFrame returns a valid frame, access the members of the NUI_SKELETON_FRAME structure.

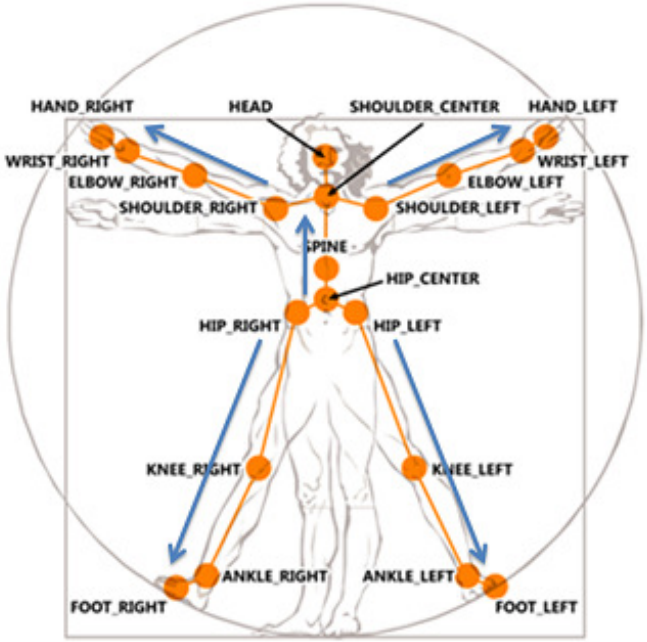

NUI_SKELETON_FRAME Structure

- Contains information about a frame of data from the skeleton pipeline.
- Syntax:

C++

```
typedef struct _NUI_SKELETON_FRAME {
    LARGE_INTEGER liTimeStamp;
    DWORD dwFrameNumber;
    DWORD dwFlags;
    Vector4 vFloorClipPlane;
    Vector4 vNormalToGravity;
    NUI_SKELETON_DATA SkeletonData[NUI_SKELETON_COUNT];
} NUI_SKELETON_FRAME;
```

NUI_SKELETON_DATA Structure

- Contains the data for one skeleton, including overall position, skeleton joint positions, and whether each skeleton joint is tracked successfully.
- Syntax:

C++

```
typedef struct _NUI_SKELETON_DATA {
    NUI_SKELETON_TRACKING_STATE eTrackingState;
    DWORD dwTrackingID;
    DWORD dwEnrollmentIndex;
    DWORD dwUserIndex;
    Vector4 Position;
    Vector4 SkeletonPositions[20];
    NUI_SKELETON_POSITION_TRACKING_STATE eSkeletonPositionTrackingState[20];
    DWORD dwQualityFlags;
} NUI_SKELETON_DATA;
```

NUI_SKELETON_DATA Structure

- Position
  - The skeleton position.
- SkeletonPositions
  - An array of Vector4 structures, each of which represents a single joint position for the skeleton. The joint positions are indexed by the values in the NUI_SKELETON_POSITION_INDEX enumeration. The x, y, and z members contain the coordinates relative to the sensor array, in meters. The w member is always unity.
- eSkeletonPositionTrackingState
  - An array of NUI_SKELETON_POSITION_TRACKING_STATE values that indicate whether the corresponding joint positions are tracked or inferred.
- The skeleton position data ranges are:
  - Values of x range from approximately -2.2 to +2.2.
  - Values of y range from approximately -1.6 to +1.6.
  - Values of z range from 0.0 to 4.0.

Tracking Modes (Seated and Default)

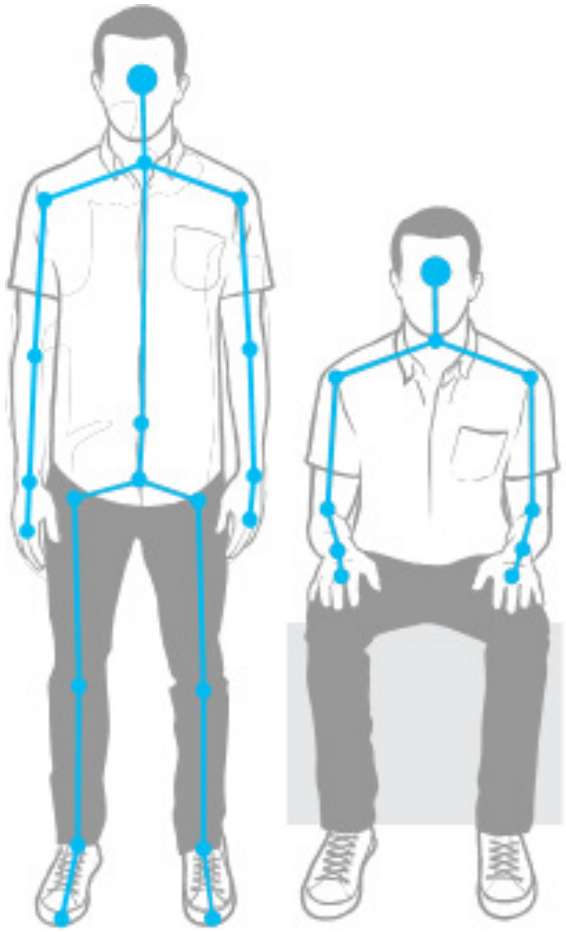

Joint Orientation

- Kinect for Windows provides joint orientation information for the skeletons tracked by the ST pipeline. The bone orientation is provided in two forms:
  - A hierarchical rotation based on a bone relationship defined on the skeleton joint structure
  - An absolute orientation in Kinect camera coordinates
- The orientation information is provided in form of quaternions and rotation matrices for use in different animation scenarios.

Joint Orientation

- Bones Hierarchy
  - Is defined a hierarchy of bones using the joints defined by the skeletal tracking system.

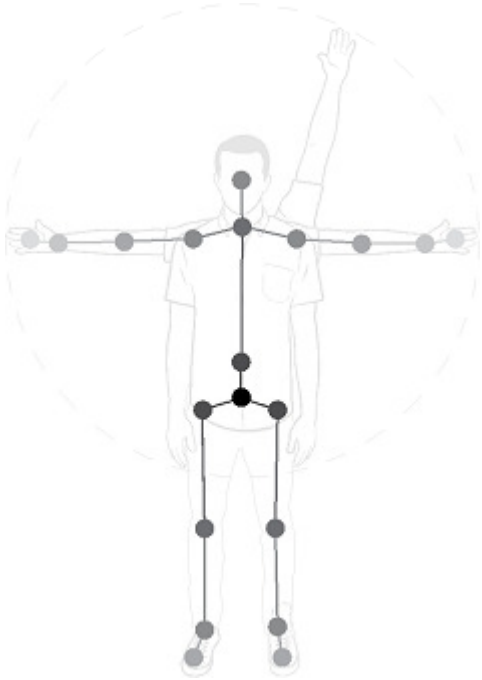

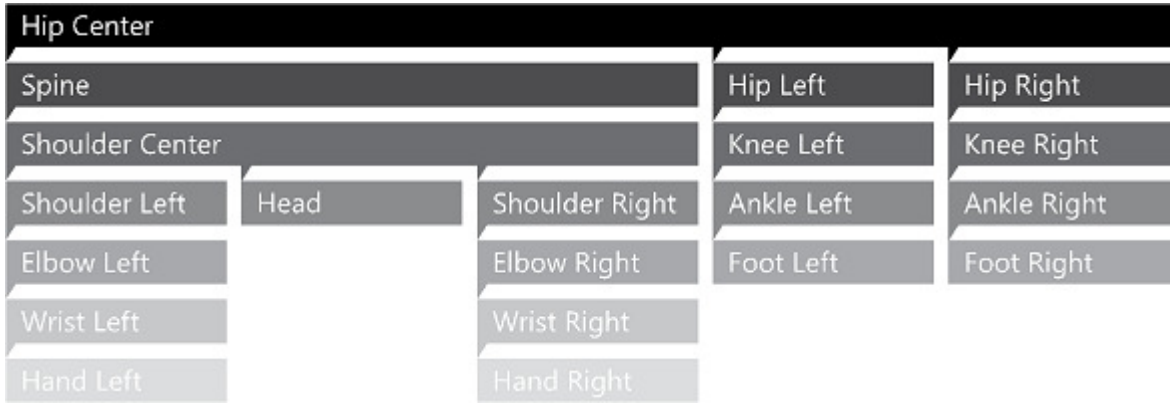

Joint Orientation

- Bones are specified by the parent and child joints that enclose the bone. For example, the Hip Left bone is enclosed by the Hip Center joint (parent) and the Hip Left joint (child).

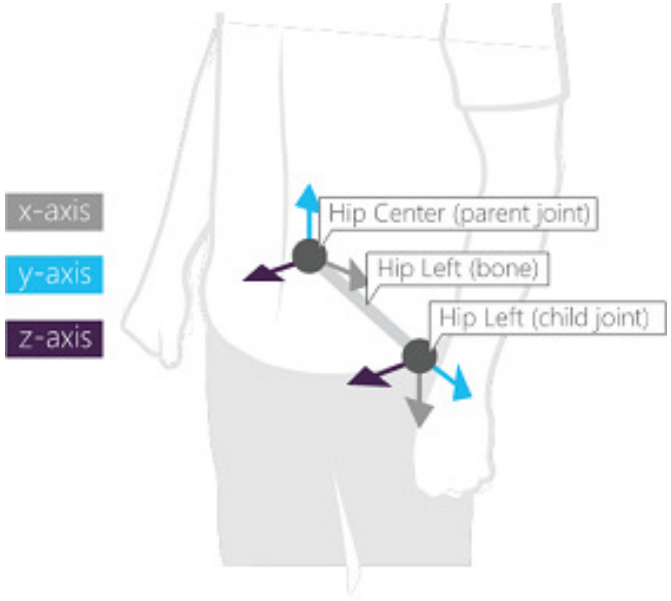

Joint Orientation

- Bone hierarchy refers to the ordering of the bones defined by the surrounding joints; bones are not explicitly defined as structures in the APIs. Bone rotation is stored in a bone's child joint. For example, the rotation of the left hip bone is stored in the Hip Left joint.

Hierarchical Rotation

- Hierarchical rotation provides the amount of rotation in 3D space from the parent bone to the child. This information tells us how much we need to rotate in 3D space the direction of the bone relative to the parent.

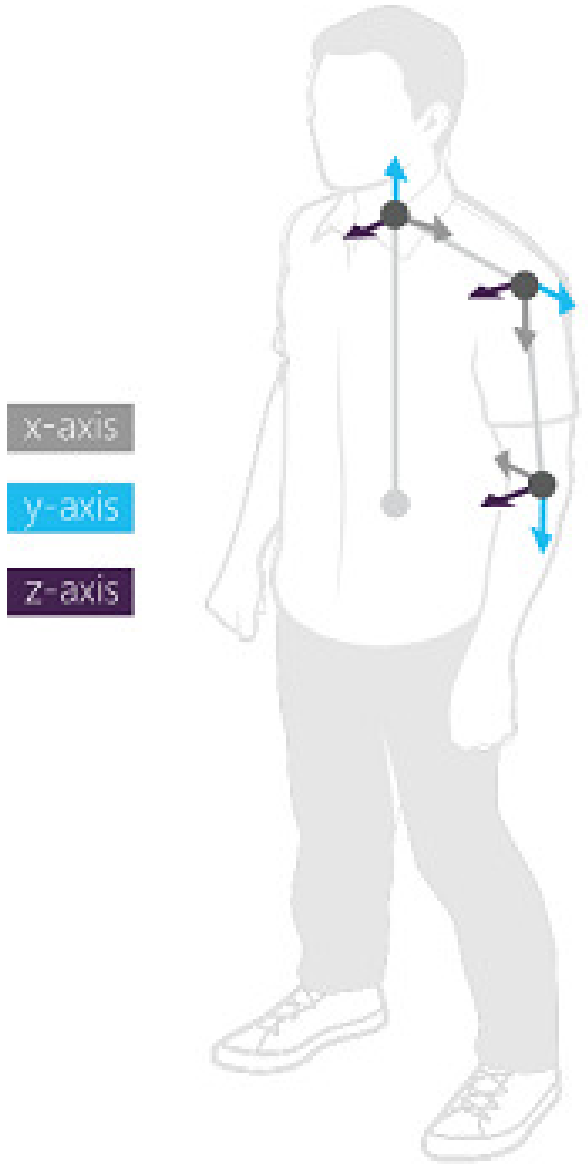

Skeleton Tracking With Multiple Kinect Sensors

- The Kinect for Windows SDK supports up to 4 simultaneous Kinect for Windows Sensors for skeleton tracking. This support is available through both the native and managed versions of the Skeleton Tracking API.

Kinect for Windows v2 hardware

- Cam HD
- Expanded field of view
- Improved skeletal tracking
- New Active IR

Applications

Rehabilitation

Iteration with people

Education

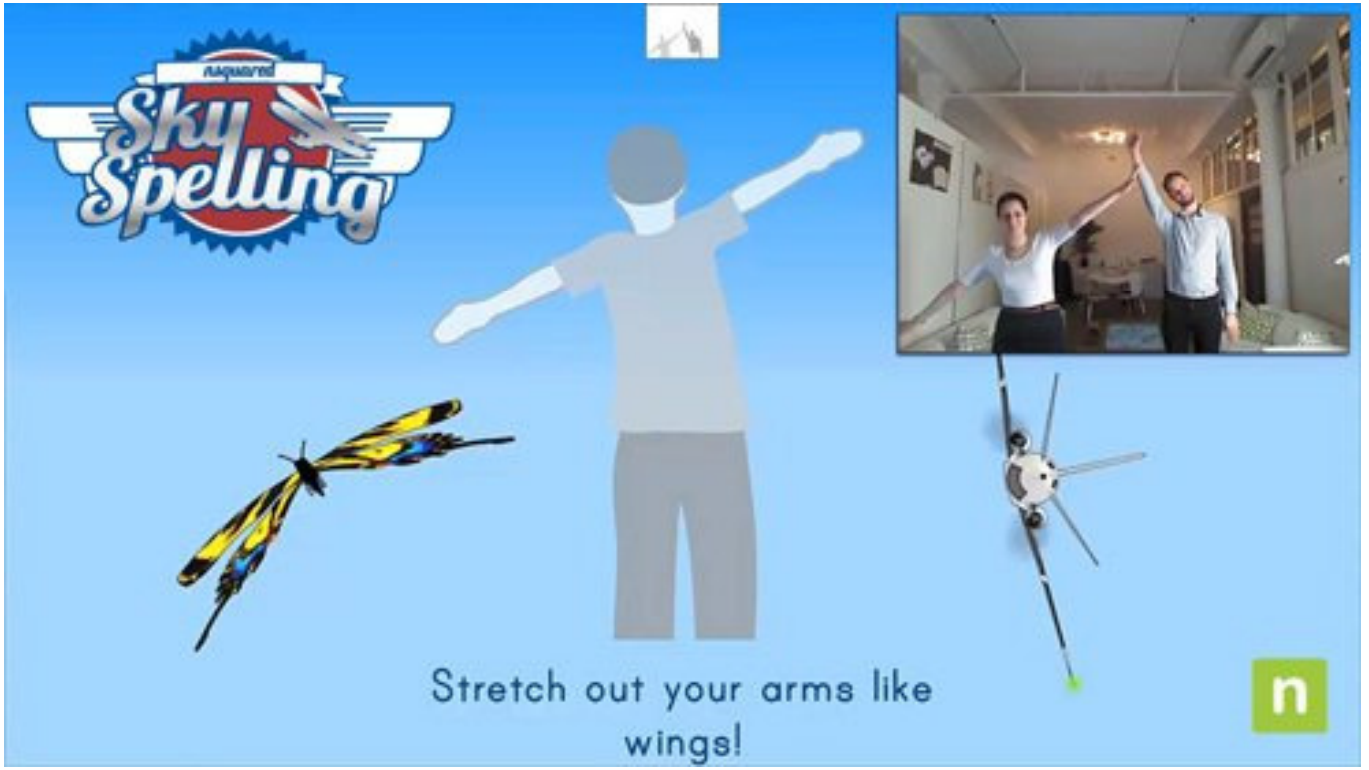

Physical Therapy with Kinect

BoxingBots: Kinect-Driven Pneumatic Boxing Robots