Enabling Fine-Grained Permissions for Augmented Reality Applications With Recognizers

Suman Jana[1*], David Molnar[2],
Alexander Moshchuk[2], Alan Dunn[1*], Benjamin Livshits[2],
Helen J. Wang[2], and Eyal Ofek[2]

[1]University of Texas at Austin
[2]Microsoft Research
* Work done during internship at Microsoft Research

USENIX Security 2013

Augmented Reality (AR) : the new frontier!

SoundWalk app on Layar AR browser

http://www.layar.com/layers/clicmobiletestlayar/

Augmented Reality (AR) : the new frontier!

EA Sports Active 2

http://www.ea.com/ea-sports-active-2

Augmented Reality (AR) : the new frontier!

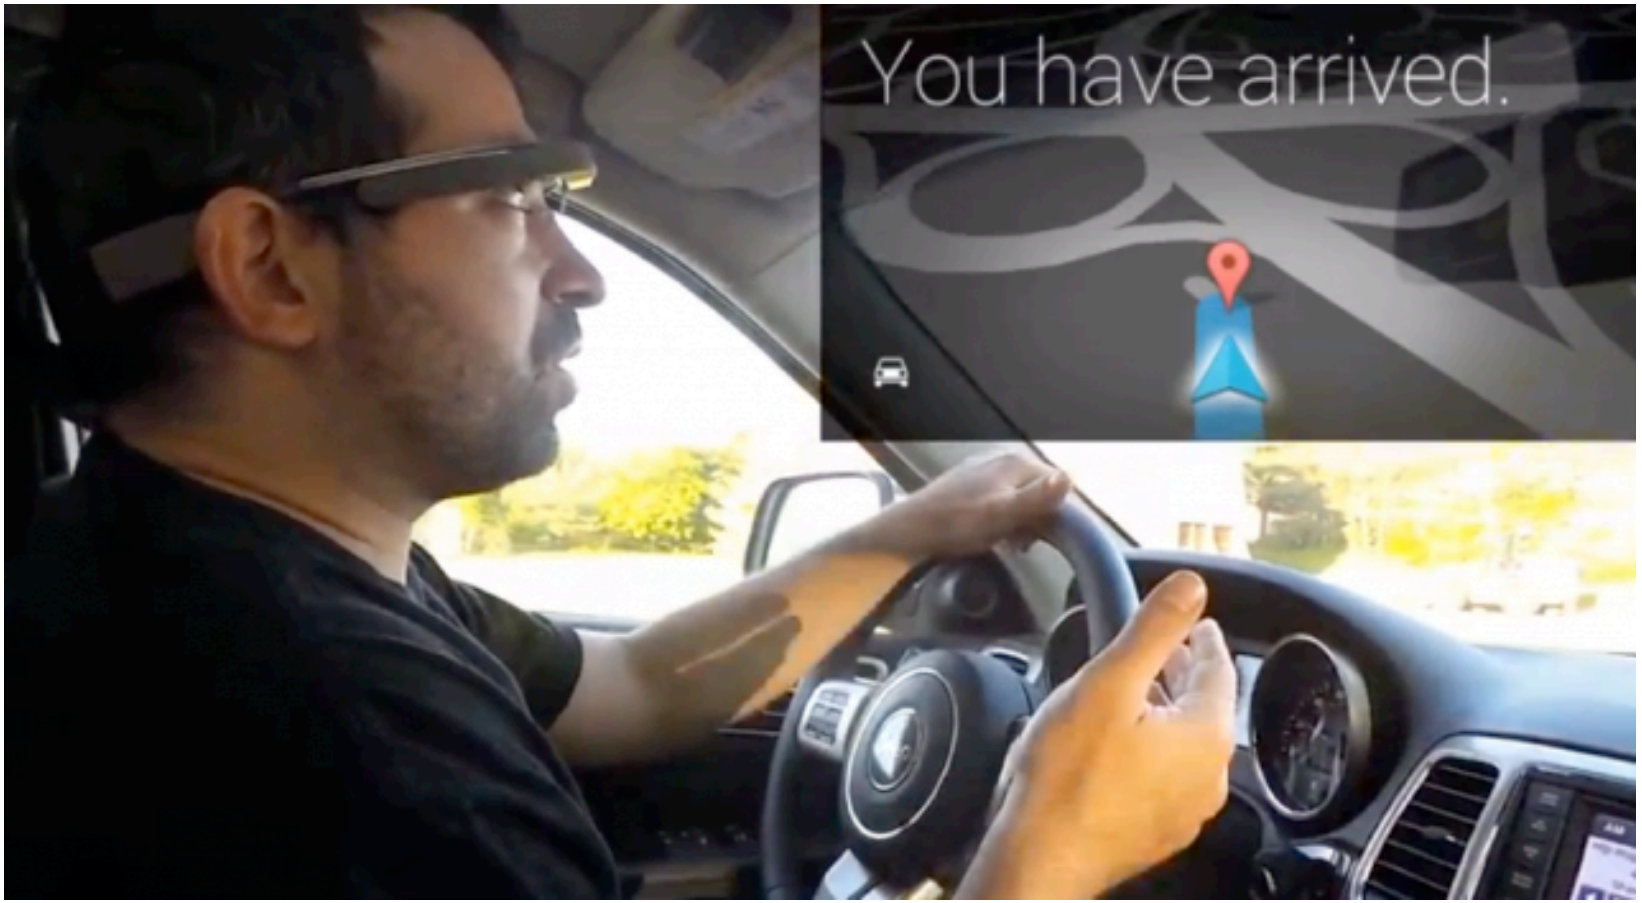

Google Glass navigation

AR application pipeline

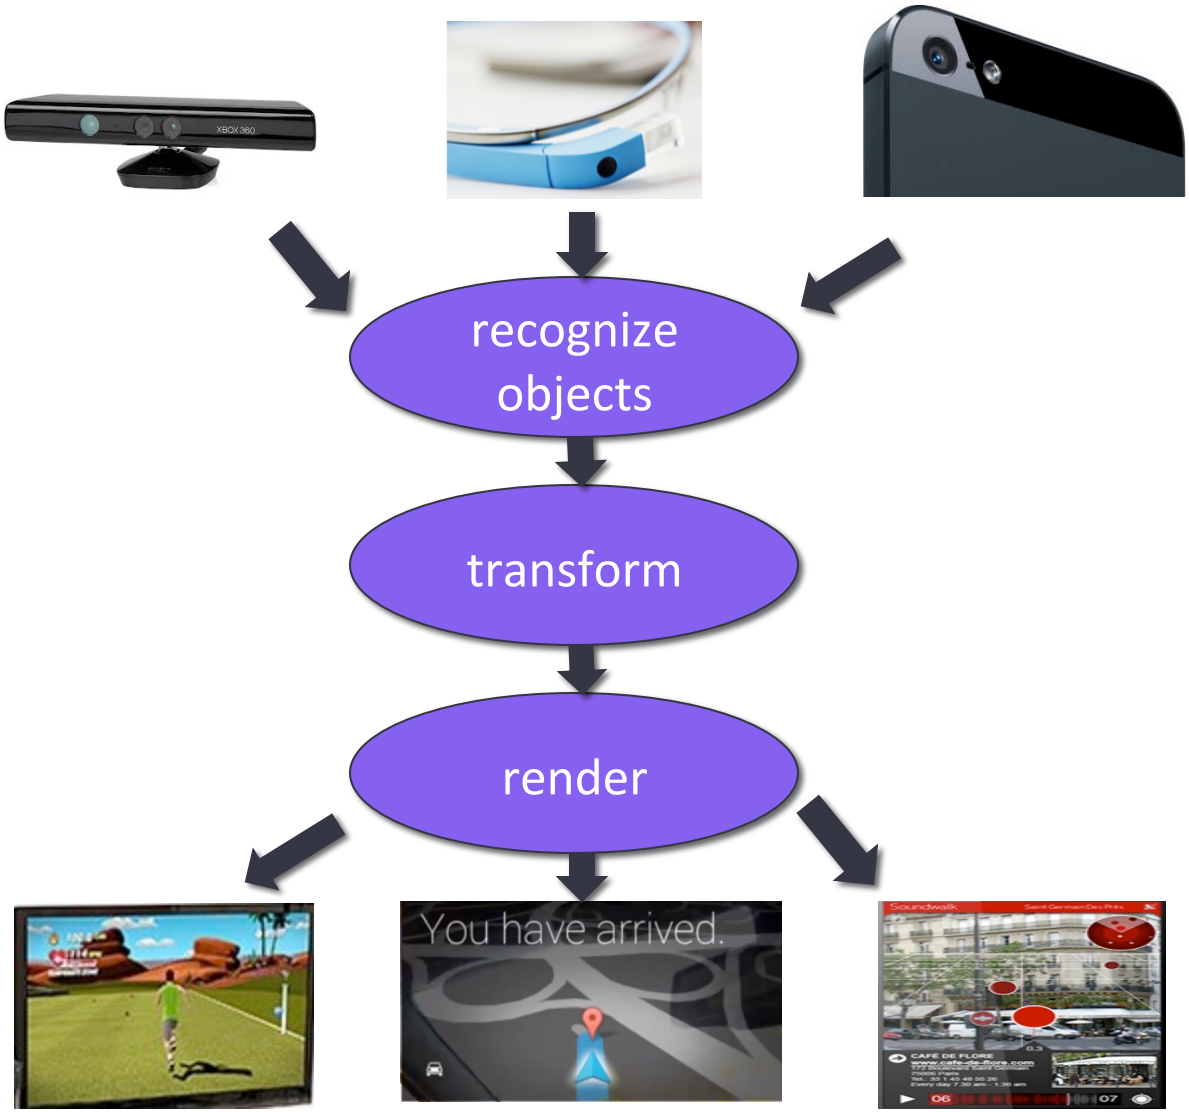

AR application pipeline

recognize objects

transform

render

topic of today's talk

current AR applications do it all by themselves!

all these stages need platform support (HotOS'13)

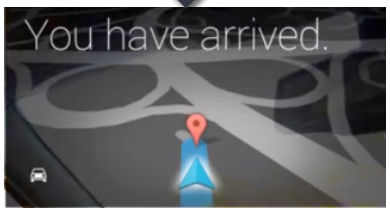

Privacy concern: unrestricted access

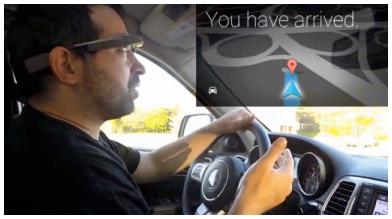

all apps see this

Privacy concern: unrestricted access

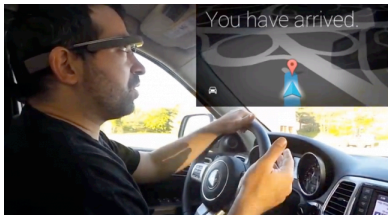

all apps see this

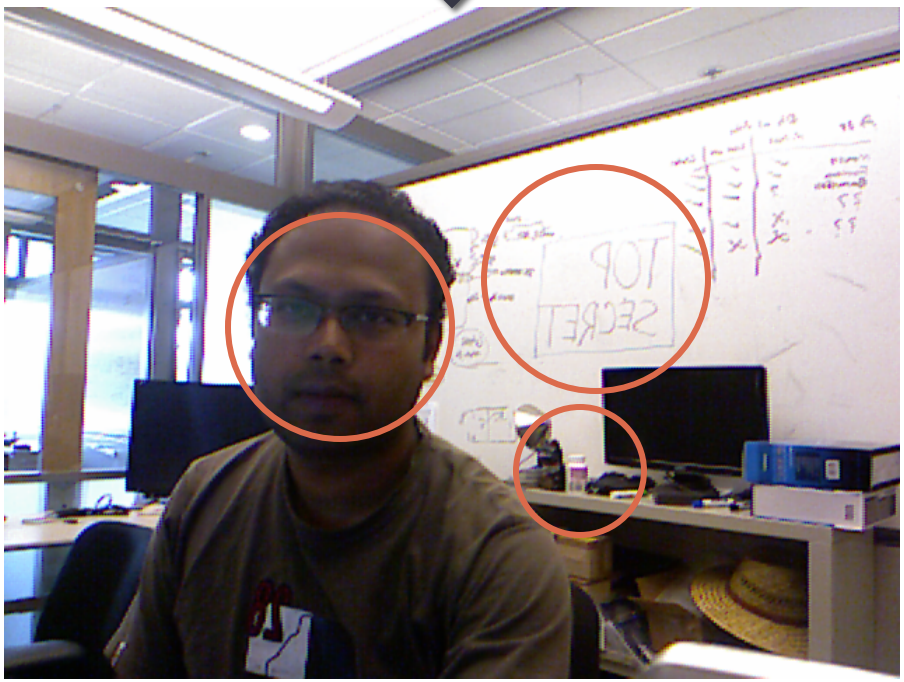

tons of sensitive information!

Functionality concern: one app at a time

AR tour guide

AR navigator

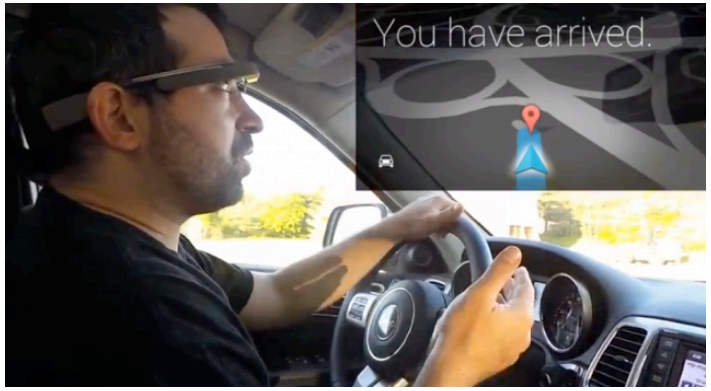

can’t run concurrently

Scalability concern: must scale to multiple concurrent applications

What can we do with today's abstraction?

a lot of repetitive work across apps

2x overhead for 2 Kinect Apps

app 1

recognize objects → transform → render

app 2

recognize objects → transform → render

hard to maintain usable frame rate

RECOGNIZERS: A NEW ABSTRACTION

Introducing recognizers

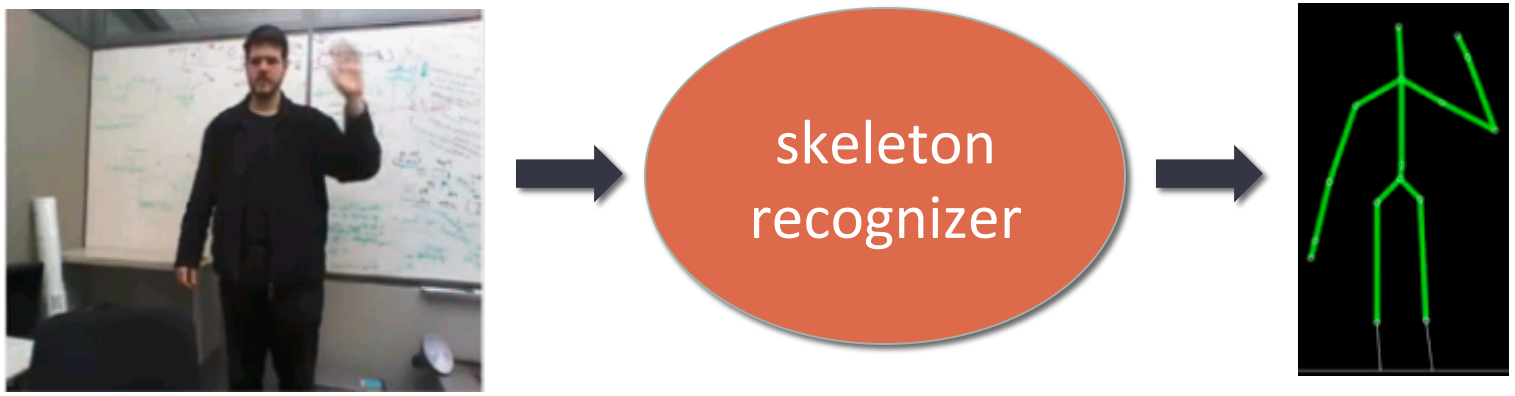

- A new platform abstraction to recognize real-world objects – generates events that apps can subscribe to
- Fine-grained control over detected objects
- Can enforce least privilege – each app must obtain permissions to access each recognizer

Scalability with multiple recognizers

Explicit data flow

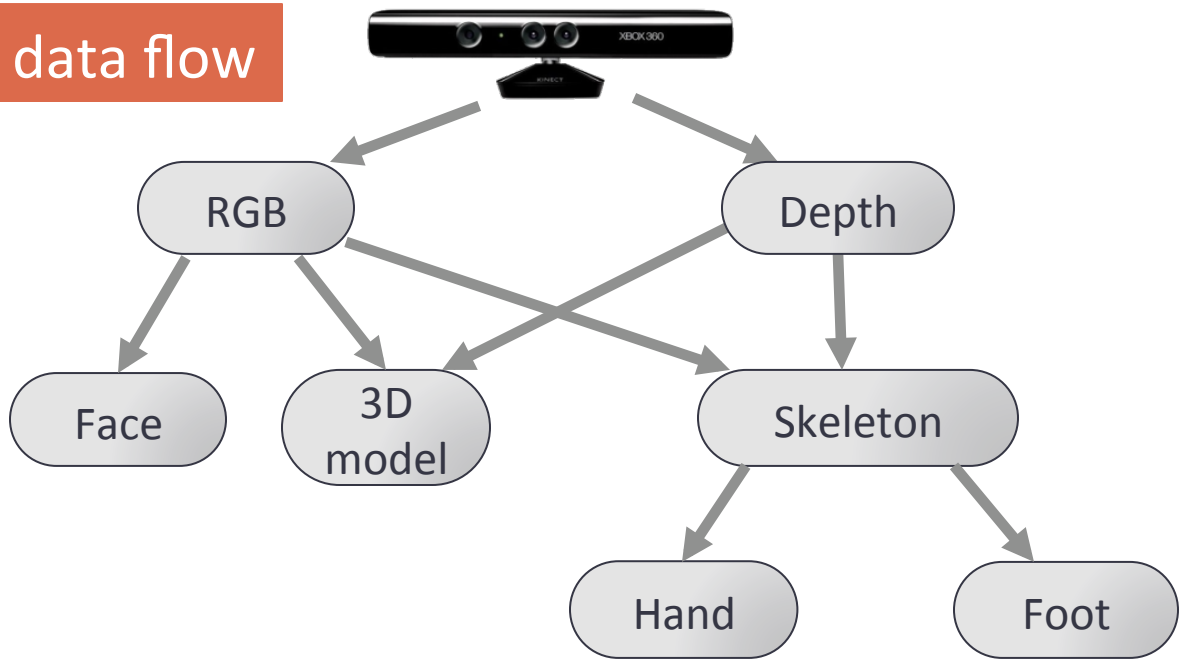

sample directed acyclic graph of recognizers

Benefits of explicit dataflow

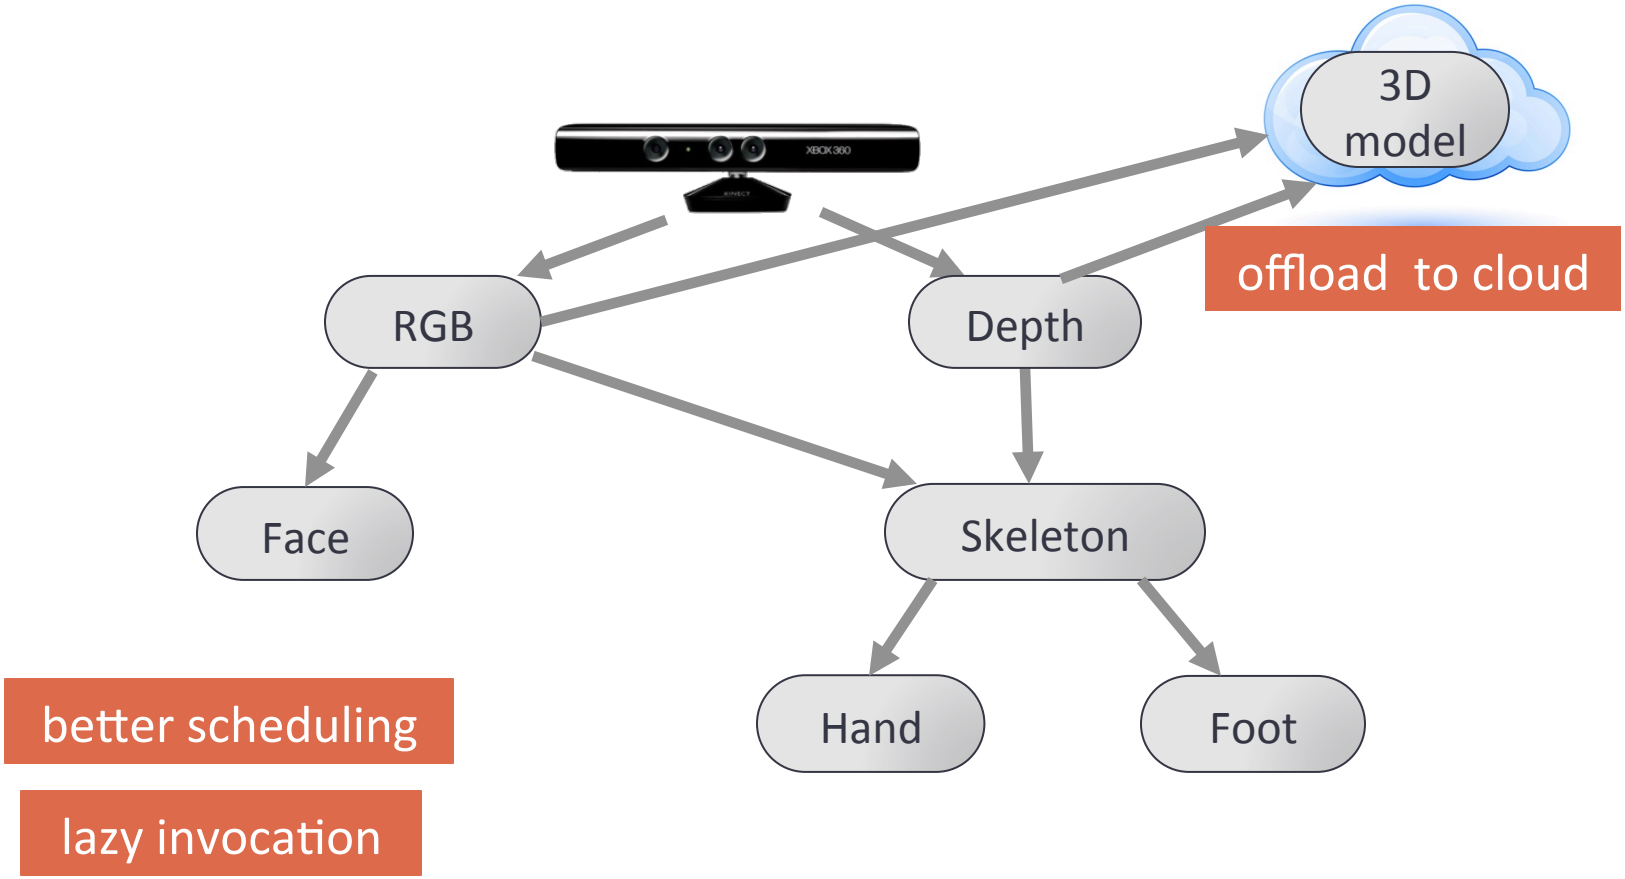

How do recognizers help?

Concerns	Role of recognizers
Privacy concern	Recognizer based permissions allow enforcing least privilege

How do recognizers help?

Concerns	Role of recognizers
Privacy concern	Recognizer based permissions allow enforcing least privilege
Scalability concern	Recognizers allow computation sharing across apps
	Offload individual recognizers

Explaining recognizers: privacy goggles

- Visual way to explain information given by each recognizer to an application

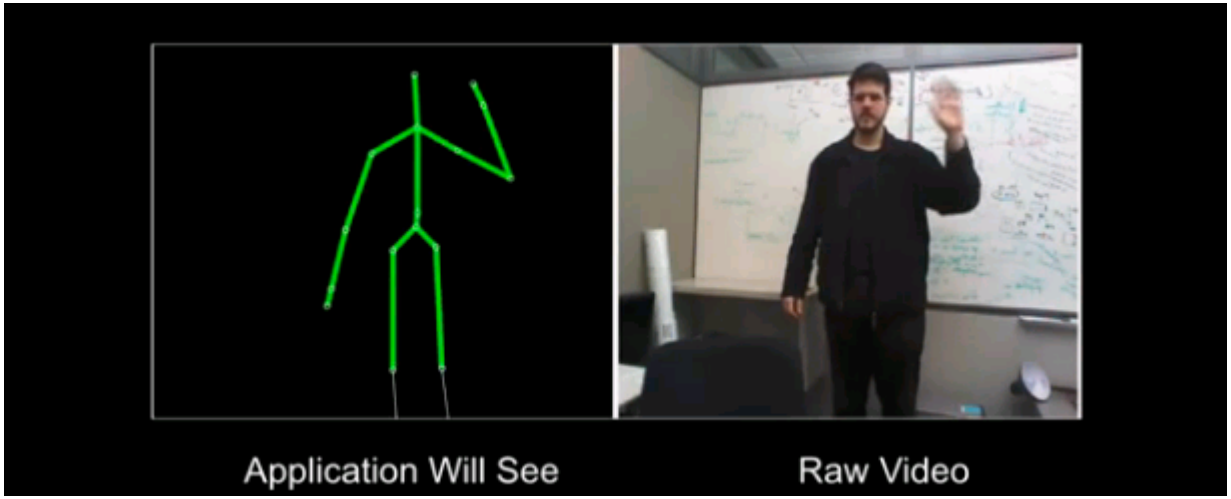

Privacy goggles for skeleton recognizer

Recognizer-based AR architecture

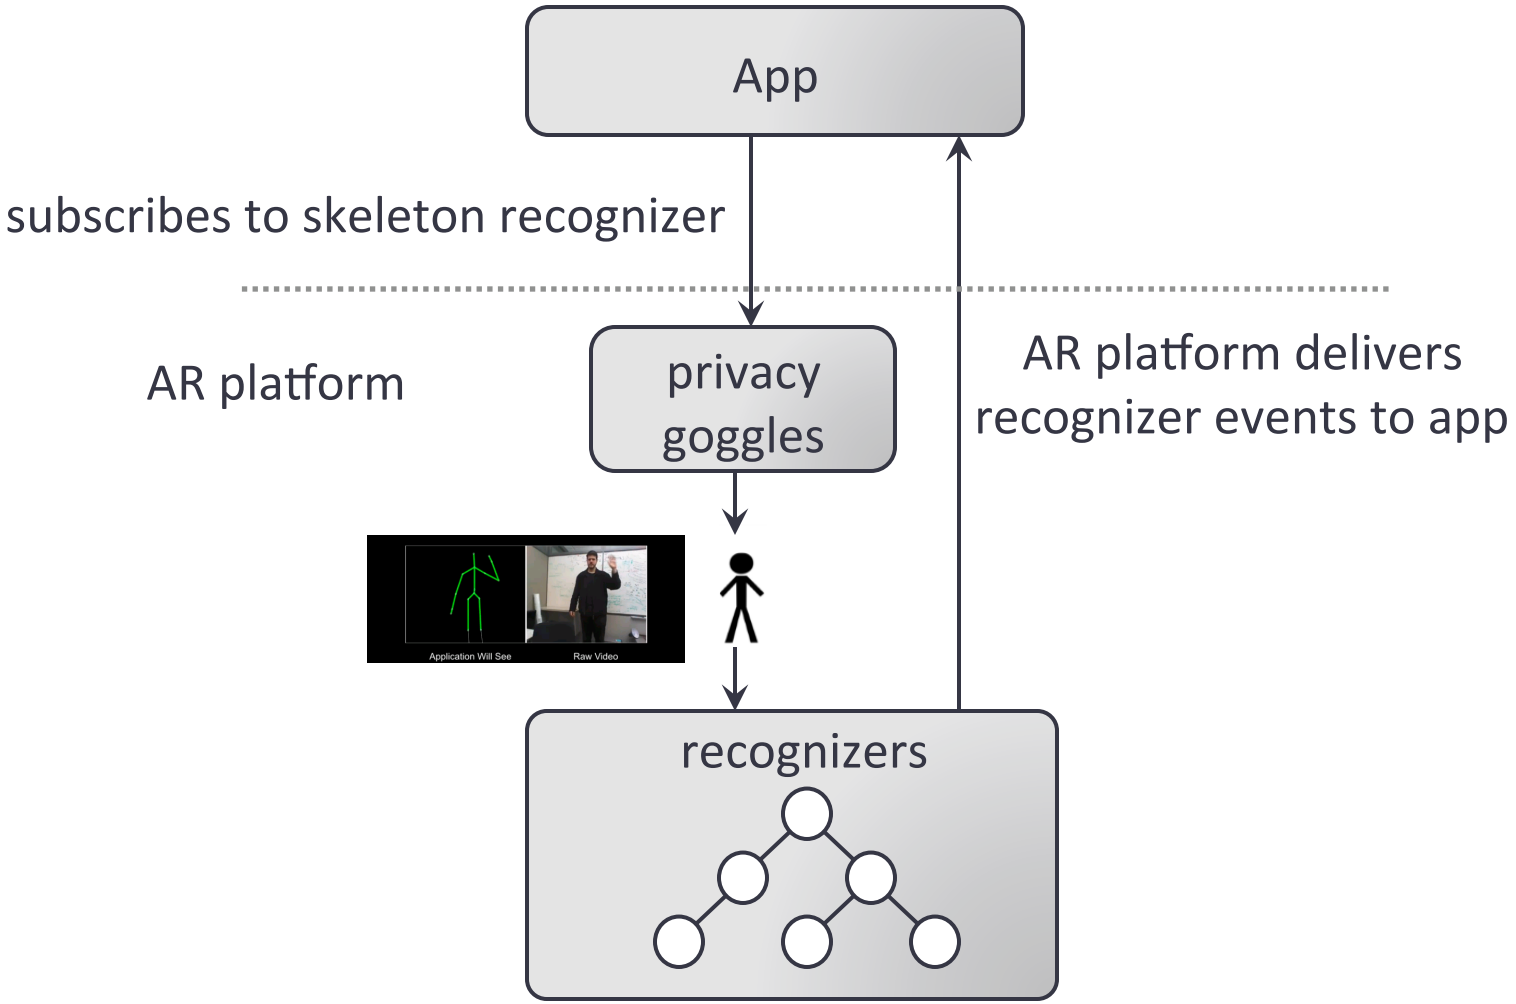

IMPLEMENTATION

Example applications

Application	What it does	Recognizers
Hand cursor	Control cursor with hand movements 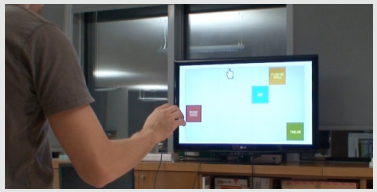	Skeleton
Facial movement detector	Visualize tracked faces 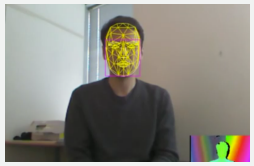	Face detector
Room Scanner	Find flat surfaces 	3D Model, Depth

EVALUATION

Evaluation criteria

- Privacy: How many applications need access to raw video and sensor data?
- Scalability
  - Performance of concurrent applications
  - Performance of outsourced AR computation
- Usability
  - Is the information released less sensitive than raw data?
  - Do users understand privacy goggles?

Few apps need raw data

Only 4 recognizers cover ~90% of shipping Xbox apps

Recognizer	% Apps
Skeleton	94.3%
Person Texture (PT)	25.3%
Voice Commands (VC)	3.44%
Hand Position (HP)	5.74%
Video Clip	3.4%
Picture Snap	1.1%
Voice Intensity	1.1%
Voice Modulation	1.1%
Speaker Recognition	1.1%
Sound Recognition	1.1%
Basketball Tracking	1.1%
Skeleton+PT+VC	82.75%
Skeleton+PT+VC+HP	89.65%

recognizers used by 87 shipping Xbox Applications

Released information less sensitive

10 surveys with 50 respondents each about recognizer output

Consider the two pictures below. Which picture contains "more sensitive" information?

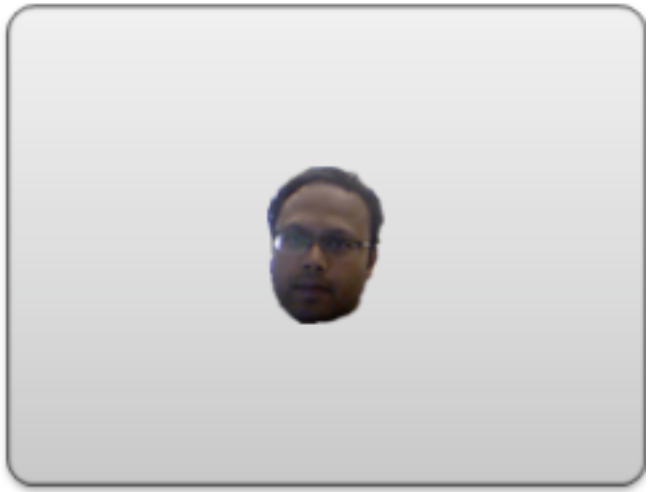

86% of the users said the left one is more sensitive

Privacy goggles communicate

- 152 respondents
- 80% identified that the app could see body position
- 47% identified that the app could see hand positions

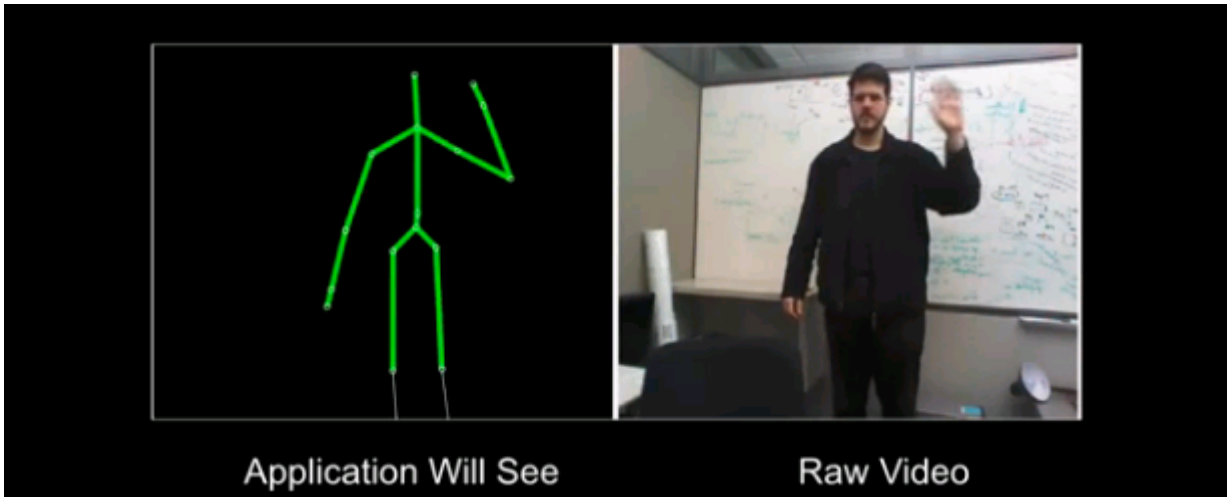

Sharing recognizers works!

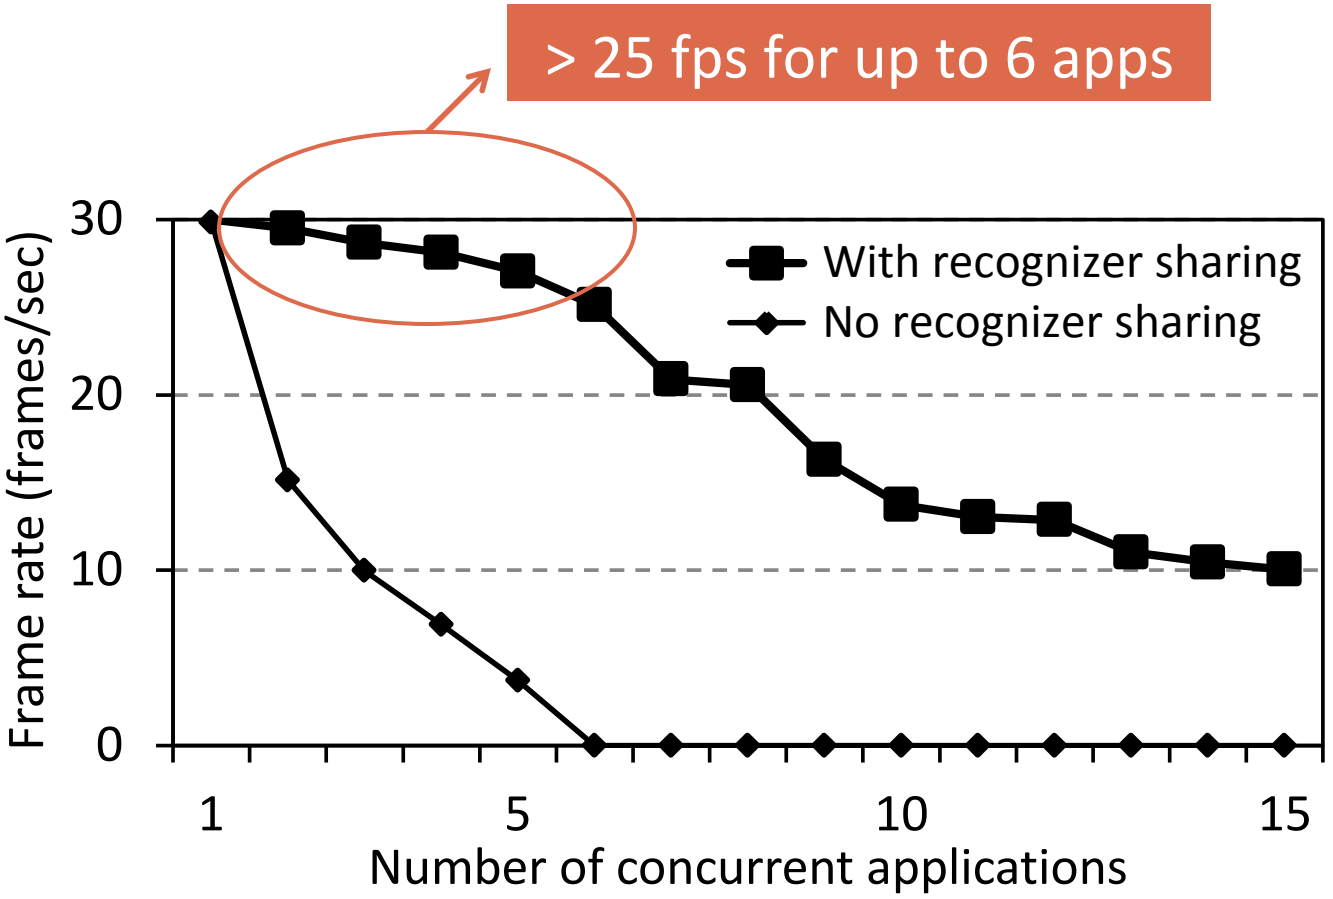

Recognizer offloading

Kinect fusion recognizer

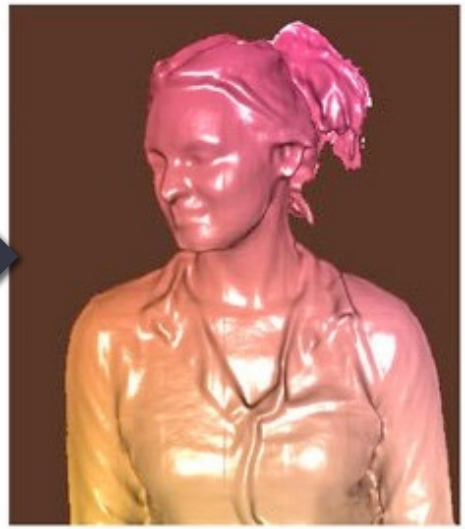

needs an extremely powerful GPU to run

doesn't work

4.46 fps

Recognizer offloading

Kinect fusion recognizer

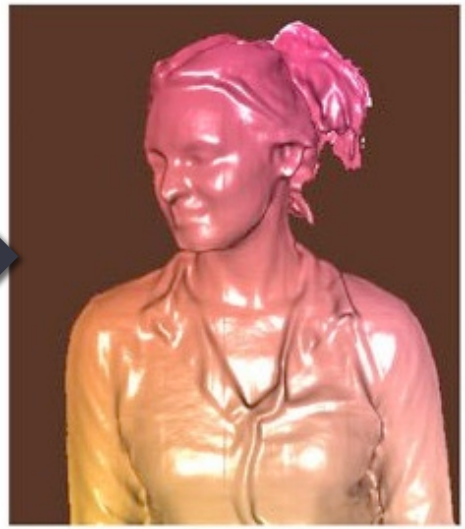

needs an extremely powerful GPU to run

4.17 fps

offloaded recognizer

4.46 fps

Minimizing recognizer false positives

- Recognizers are not perfect (yet) - false positives can lead to information leakage

AR platforms can apply simple heuristics like combining multiple recognizers to decrease false positives

OpenCV face recognizer

OpenCV face recognizer & Kinect depth filter

Summary

- New AR paradigm needs new platform abstractions
- Recognizers
  - Help in enforcing least privilege
  - Allow sharing of computation across apps
  - Allow efficient offloading of heavyweight computation
  - False positives in recognizers can be dealt using heuristics at the platform level
- Privacy goggles - visual permission management

Future work

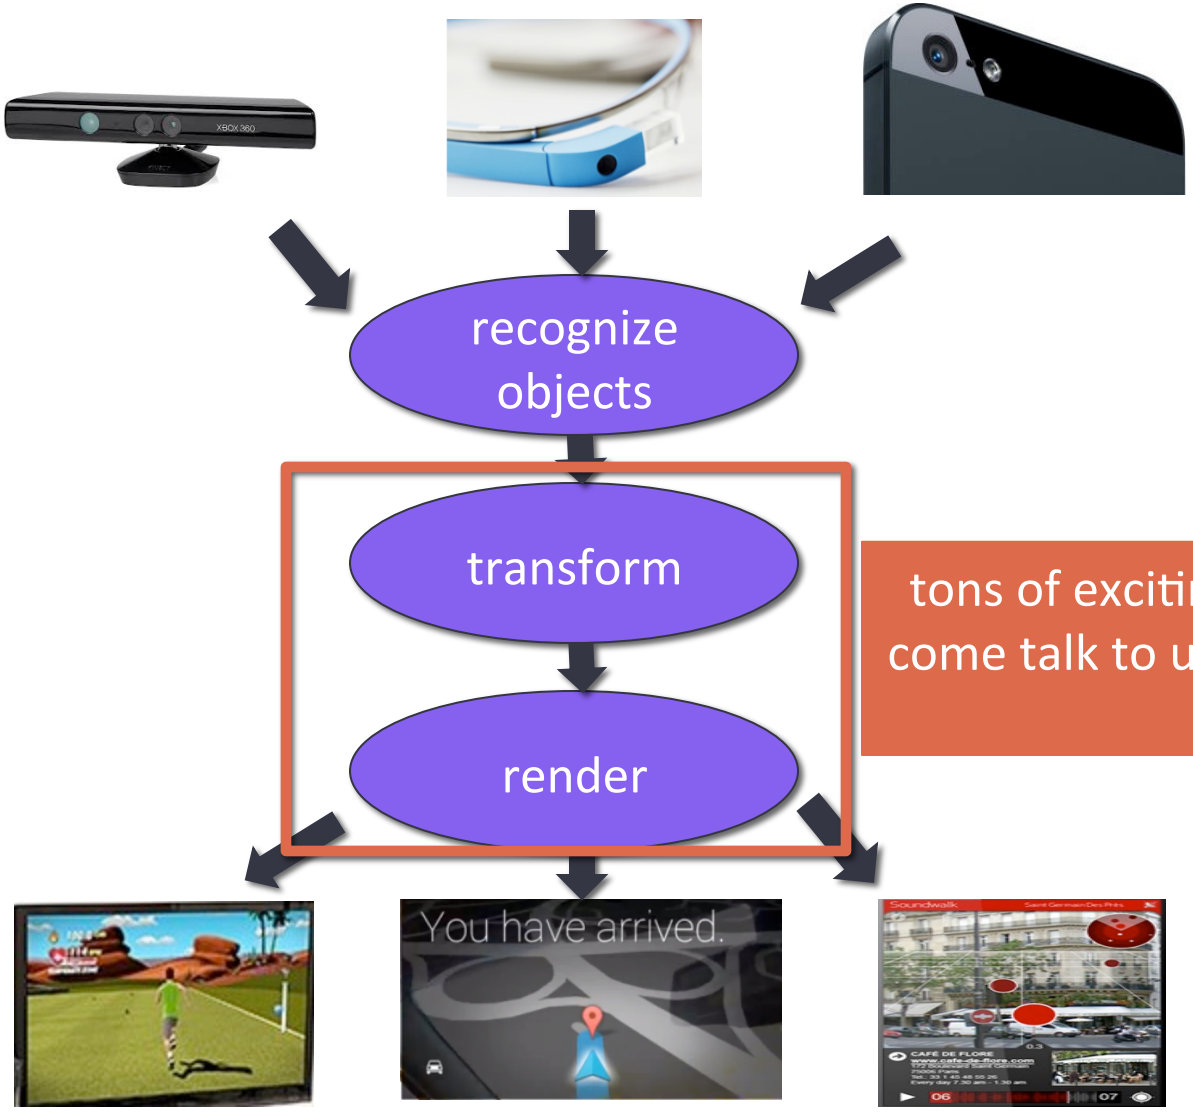

tons of exciting new stuff!
come talk to us if interested

Thanks!