VR/AR Association White Paper

Introduction to the AR Cloud with Use Case Examples

Author:

Alex Chuang, Shape Immersive

Amy LaMeyer, EnteringVR

Colin Steinmann, Bent Image Lab/youAR

Gabriel Rene, VERSES

Mikko Karvonen, Immersal

Sam Beder, Ubiquity6

Steven Swanson, VERSES

Contributors:

Matt Miesnieks, 6d.ai

Kris Kolo, VRARA

Table of Contents

1 Introduction to the AR Cloud

The world is moving towards a fundamental shift where our physical reality will soon blend with a virtual one. This idea opens up an entirely new frontier in which our experiences and our realities will be extended in ways we could have never imagined.

In this near future, the possibilities for Augmented Reality (AR) are endless. Brands can attract and engage customers with more immersive and interactive experiences not bounded by physical constraints. Employees can learn how to operate equipments more effectively in complex assembly lines, reducing cost and risks for businesses. Students can visualize complicated diagrams in 3D, improving academic performance. Consumer products, instruction manuals and textbooks are just a small fraction of static objects that can be brought to life.

Up until now, AR experiences have been rudimentary and siloed primarily because they were hard to develop, hard to distribute and had no real demand. Most so-called AR experiences have been merely a simple 2D digital overlay on the real world with no real connection between the virtual content and our physical world. For example, in the original Pokémon Go, the AR characters do not understand the spatial context of the surrounding area.

In order to consume context-aware AR content in the physical world, it is necessary to understand the precise location and orientation of the viewer's device. As pervasive and useful as GPS is, its global average user range error (URE) is ≤7.8 m. This is far too inaccurate for virtual content to align with the physical space, particularly in confined spaces. Due to this inaccuracy, two people in the same space will end up with two different frames of reference, making synchronized multi-user experience impossible.

For mass adoption of AR to occur, content must persist in the real world across space, time and devices. As Ori Inbar, Partner of Super Ventures, describes, "persistence means I can create an AR experience in a physical space today, come back tomorrow and interact with it. Another user can also collaborate with it on a different device." For example, an artist can place their 3D virtual art in a plaza for the community to interact with it, regardless of which device they use. The 3D virtual art will "live" in that space as if it's really there and will not disappear between different app sessions. Multi-user, occlusion are two additional functions that are key to augmented reality adoption. To enable these abilities and a streamlined experience, the "AR Cloud" is needed.

1 1.1 Definition(s) of AR Cloud

According to Matt Miesnieks, CEO of 6D.ai, the AR Cloud is a “machine readable 1:1 scale model of the real world” or others have called it “a real-time spatial map of the world”. Think of the AR Cloud as a parallel universe of our world where objects, buildings, nature have been digitized. Inbar calls AR Cloud “the single most important software infrastructure in computing, far more valuable than Facebook’s social graph or Google’s PageRank index”.

The VR/AR Association AR Cloud Committee (http://www.thevrara.com/ar-cloud) defines the AR Cloud as the following:

The AR Cloud is a continuously updated collection of machine-readable datasets, primarily sparse or dense point clouds plus a feature descriptor and other meta-data for each point or groups of points. AR Cloud data is any dataset which can aid in accurately determining the position and pose of AR-enabled devices, semantically understanding the scene and positioning digital content in physical spaces (this can include polygonal geometry, visual descriptors, images, video, ‘simple geographic features’, etc).

Because of the ever-changing nature of real-world spaces, the AR Cloud must constantly be updated with new data to enable AR devices to accurately localize within these dynamically changing environments. AR Cloud data is crowdsourced from an array of technologies (sparse slam maps, point clouds, occlusion meshes, photorealistic textures, BIM data) and versioned to reflect multiple states including time of day, environment changes, weather or lighting changes.

Through platforms like Ubiquity6, 6d.ai, and others, the AR Cloud enables:

- Persistent AR, indoors and outdoors
- Multiplayer AR
- Occlusions of static or semi-static (parked cars, tree foliage) objects
- Neural-network training data repositories
- Content meta-data (e.g. coordinates of persistent content) hosting
- Content state management
- Semantic scene understanding
- Presentation of real-world 3D data in various forms (sparse slam maps, point clouds, occlusion meshes, photorealistic textures, BIM data)

Of note, Spatial Web is a broader term consisting of several technologies among which are AR Cloud solutions. The Spatial Web is the convergence of IoT, Self Sovereign Identity, Artificial Intelligence, Edge Computing, Spatial Browser, Digital Ledger Tech and others.

1 1.2 Tech Giants Are Investing in AR

Apple's CEO Tim Cook said, "I see AR as being profound. It has the ability to amplify human performance instead of isolating humans, and so I'm a huge believer in AR. We put a lot of energy in AR, and it's moving very fast."

He's right, things are moving fast. In 2017, five major AR development platforms were announced. With the release of ARKit and ARCore, Apple and Google are a string position in aiding AR through the novelty phase. With ARCore and ARKit, developers can currently make applications that utilize (1) The phone's 6 degree of freedom pose, with new coordinates each session; (2) a partial & small ground plane; and (3) a simple average of the scene lighting.

Joining Apple and Google are Facebook's AR Studio, Snapchat's Lens Studio and Amazon's Sumerian. When the world's most valuable and powerful tech giants take a serious leap forward in AR, one can assume that there will be a proliferation of AR applications in 2019 and beyond.

NAME	ARKit	ARCore	AR Studio	Lens Studio	Sumerian
RELEASE	June, 2017	May, 2017	April, 2017	December, 2017	November, 2017
DIFFERENTIATOR	A11, Early Adoption, Built in AR (Animojis)	Integration with other Google Products	Emphasis on Facial Augmentation	Snapcodes, Templates	Platform-agnostic, In-Browser Dev, Hosts
GOOD FOR...	Sand Box Development for Apple Apps	Sandbox Development for Android Apps	Creating Shareable, Social AR Experiences	Producing Easily-Broadcasted Social AR Experiences	Business & Enterprise Training

Image Source: https://medium.com/@istrategylabs/augmented-reality-everything-you-need-to-know-for-2018-7988ffab1f61

Benedict Evans, Partner at a16z, believes that AR could be the next fundamental platform shift; a universal interface that would replace the multitouch. Evans describes new technology tends to follow an S Curve. First, things are slow in the R&D phase as fundamental concepts are being worked out. Then, there's a period of rapid changes, innovation and feature expansion. Finally, as the market matures, things slow down and innovation becomes incremental. We've seen this cycle in the mobile era where the first announcement of iPhone in 2007 launched a whole smartphone revolution. Compared to other technology platform shifts, "AR today is where smartphones were 10 years ago and if there's any predictability of technological cycles, in three years' time, AR will be in the frenzy phase of the S curve."

1

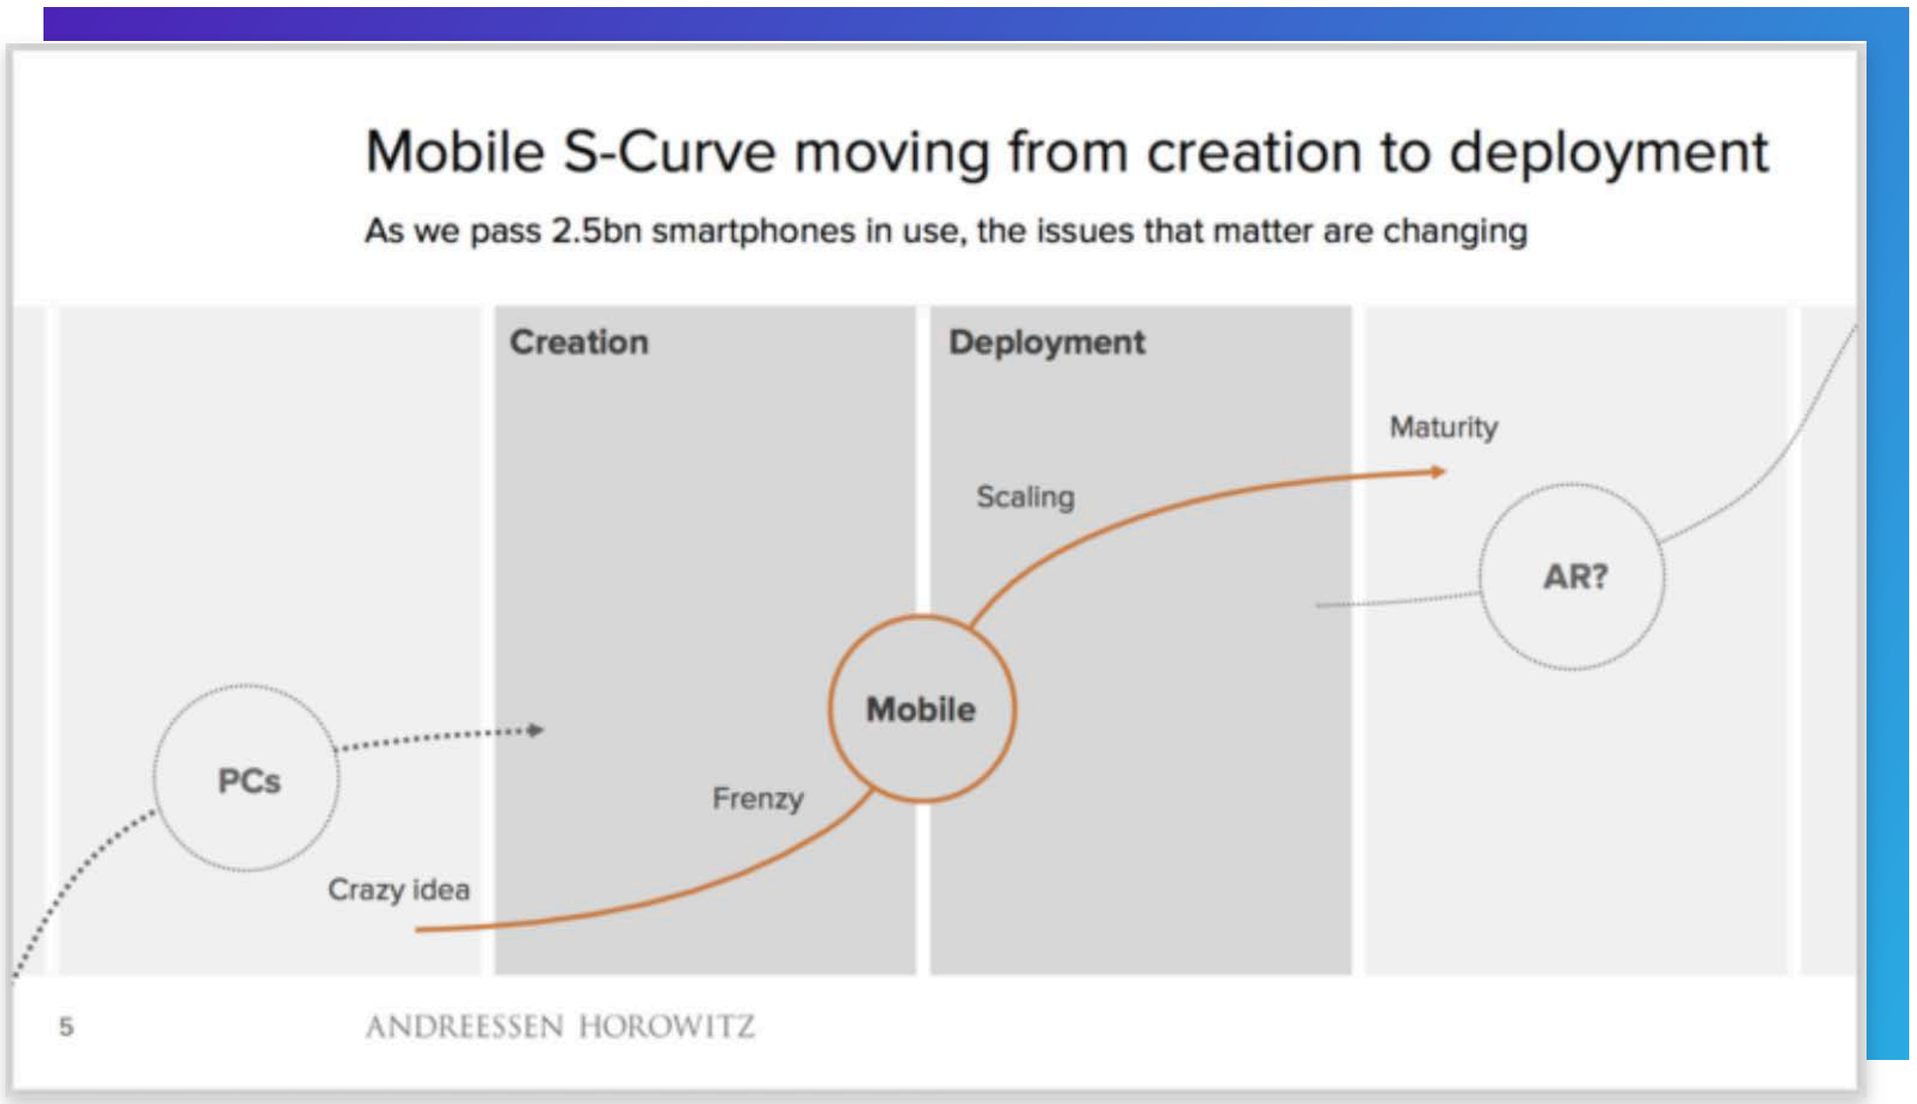

Image Source: https://www.ben-evans.com/benedictevans/2017/3/22/the-end-of-smartphone-innovation

1 1.3 Building the AR Cloud

There are three key major components of the AR Cloud that Inbar has laid out:

1. A scalable and shareable point cloud aligned with real world coordinates plus associated meta-data
2. The ability to instantly localize (align the world's soft-copy with the world itself) from anywhere and on multi devices
3. The ability to place virtual content in the world's soft-copy and interact with it in realtime, on-device and remotely

In a three-dimensional coordinate system, a point cloud is simply a set of data points defined by x, y and z coordinates and often intended to represent the external surface of an object. Point clouds are typically dense data points so they are often converted into lighter polygon meshes to represent the geometry of a surface.

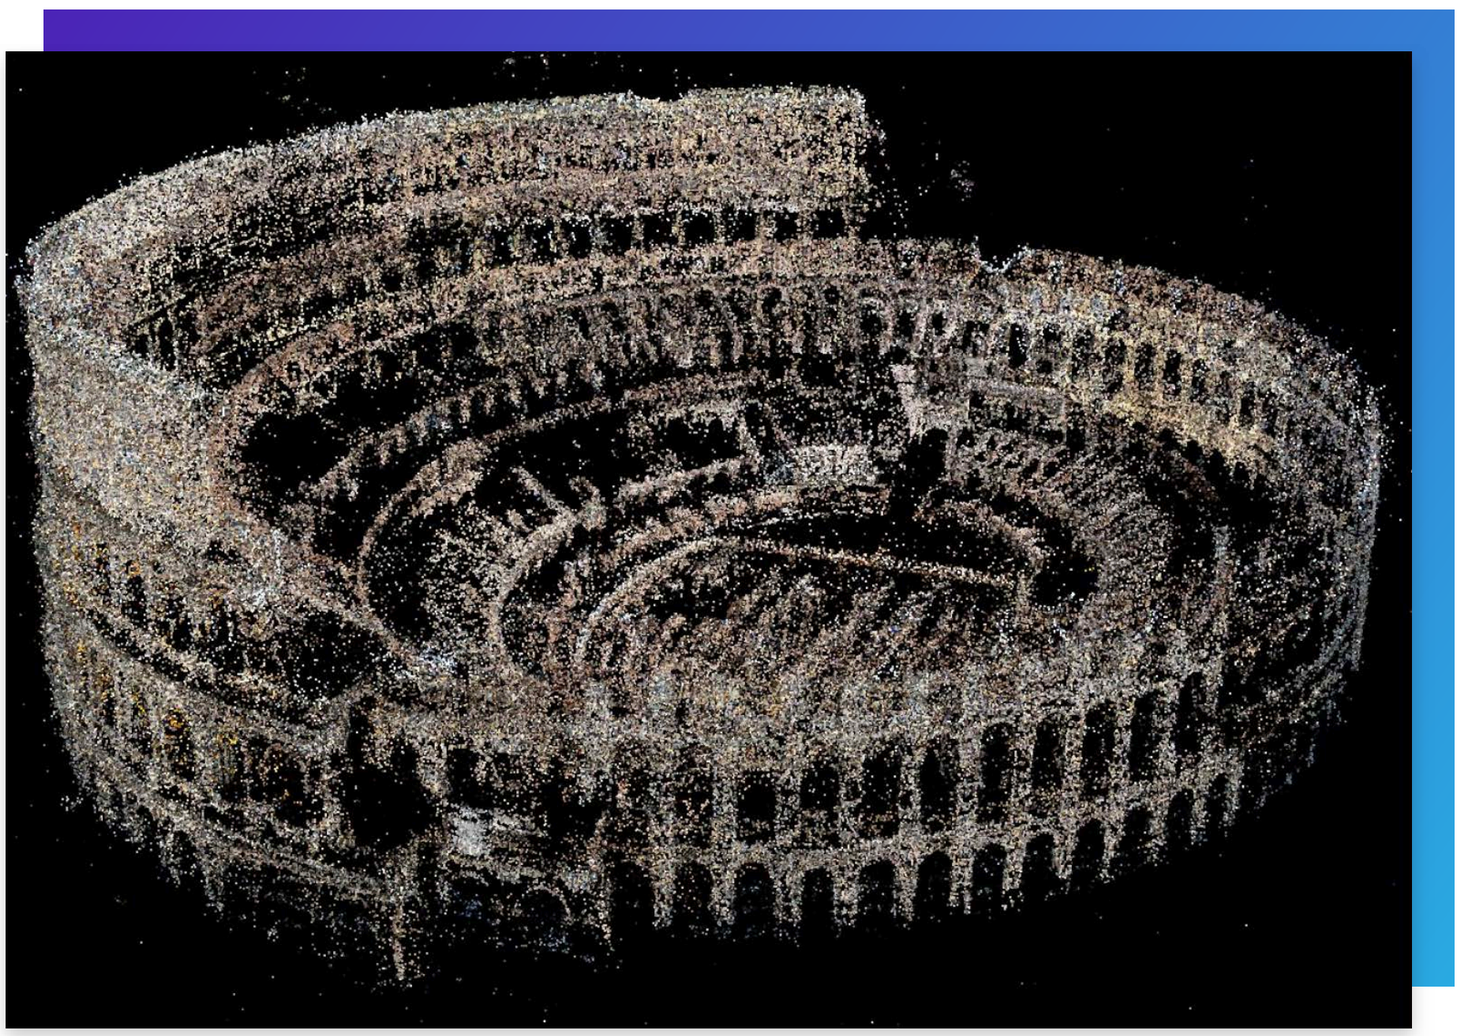

Image Source: https://www.worldviz.com/post/converting-point-clouds-into-virtual-reality-using-vizard

For the AR Cloud, the point cloud must be persistent, accessible and align with real-world coordinates. It needs to have high enough resolution to support occlusion (hiding the virtual object behind a physical object), collision (colliding a virtual object with a physical object) and extracting semantic meaning from the scene (eg segmenting objects from each other, identifying objects, determining physical properties of a surface etc).

To illustrate this, GeoSim, a Vancouver-based company, has created the most accurate and detailed 3D model of Vancouver1 by scanning Vancouver from the air and ground. They are able to fuse data collected from aerial LiDAR/photography, and ground LiDAR/photography to create a high fidelity 3D model of the city with 1 cm visual resolution and 5–10 cm spatial accuracy. This is the type of accuracy that is needed for the AR Cloud and unfortunately GPS cannot produce the same level of spatial accuracy.

1

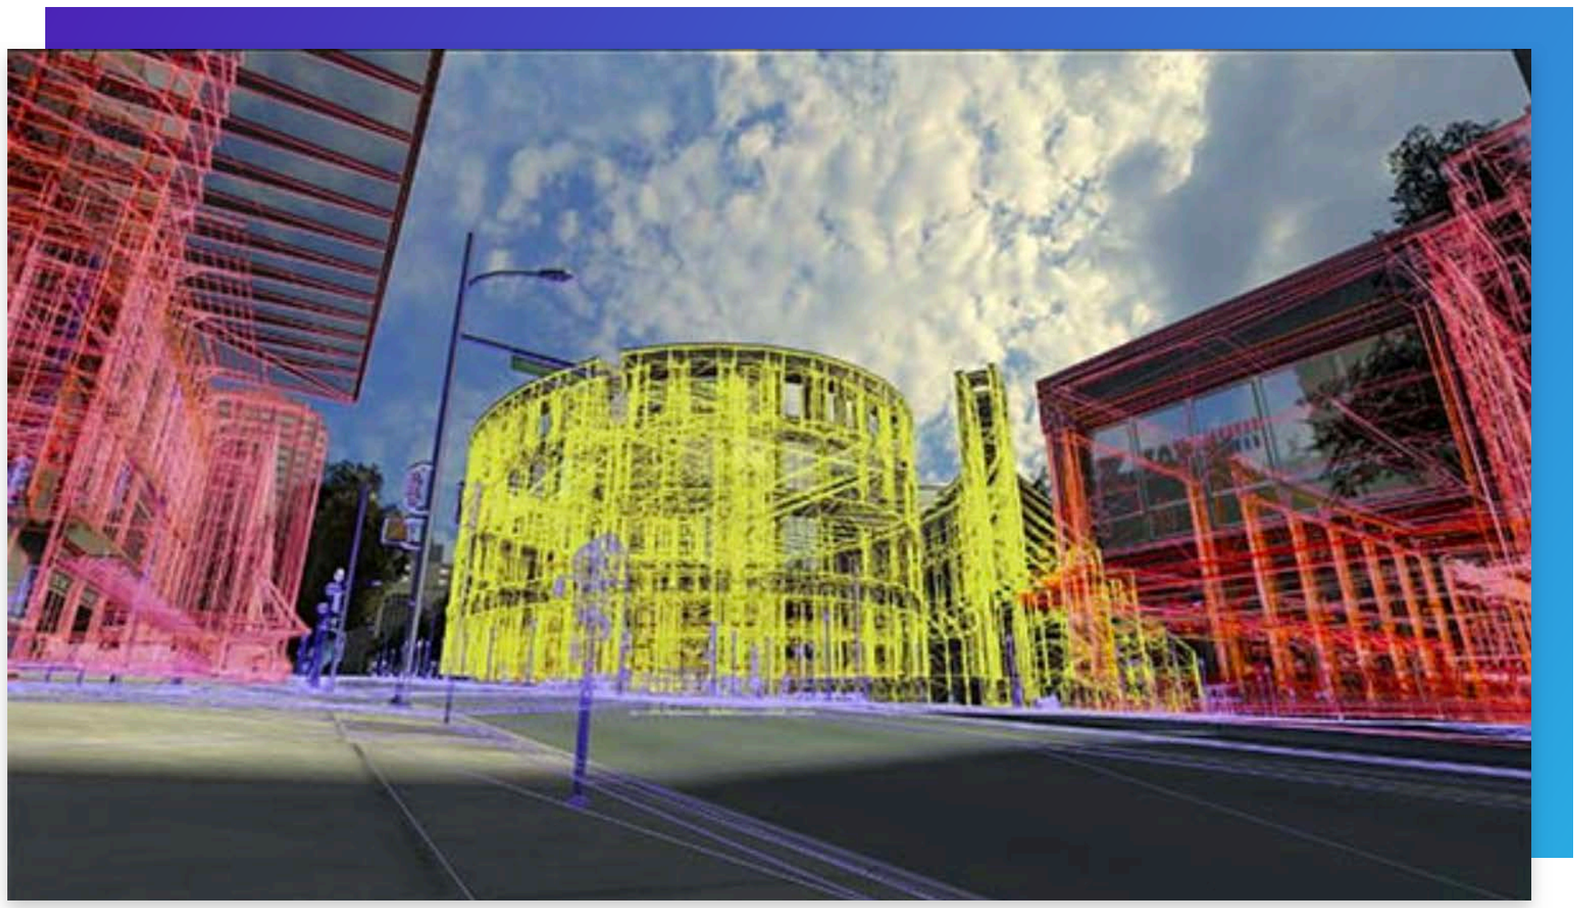

Image Source: http://geosimcities.com/

The reason why we need such a detailed point cloud of our world is because we need reference data to localize our devices. Localizing a device simply means that the camera on the device needs to understand its precise spatial coordinates relative to its environment.

With the help of the right computer vision technology and robust point cloud, the device will be able to instantly compare key feature data points with the data points in the AR Cloud to find a match. Tracking is then done by an application running on the user's device which use AR SDKs (software development kits) such as ARKit for Apple, ARCore for Android devices, or others. Apple's ARKit, for example, uses a technique called visual-inertial odometry which combines information from iOS device's motion sensing hardware with computer vision analysis of the scene to determine the position and motion of the device.

Localizing the device against a small local point cloud or a small set of point clouds has been done time and time again. However, the ideal AR Cloud localizer will be able to localize against a vast set of local point clouds from any given angle and can share the point cloud with multiple cross platform devices. Many companies around the world are currently solving this problem, making persistent and multi-user AR experiences possible. By using a smartphone's passive camera feed, they're able to reconstruct any scene in 3D and share this data with other cross platform devices. To explain it simply, if your iPhone had already scanned your living room, your friend's Android device can also access that data and enjoy the same persistent AR experience in your living room.

This brings us to another component of the AR Cloud: the ability to position and visualize virtual content registered in 3D. This opens up a new possibility for multiple users to interact with the virtual object in real-time, at various angles and remotely. Each person will be able to experience the virtual object in his or her unique point of view, just like how they would interact with a real hologram.

[1] https://vimeo.com/156995925

2 Use Cases

As a new computing platform, the possibilities for AR are endless and will impact all industries. Many of them will rely on the AR Cloud as in the examples below. Additional use cases include training, music experiences, education, travel, real estate, healthcare retail, and social messaging.

2 2.1 Gaming Niantic

Launched in June 2016, Pokemon Go by Niantic was the first augmented reality game to have mass success. Over the years as the AR cloud features grew (occlusion, multiplayer), so has the associated revenue. Reported by VentureBeat, the location-based AR game made $68.1 million in January 2019, an 84 percent increase over January 2018. Total revenue is over $2.3 billion.

Niantic has larger sights than gaming. They are building what they describe as a planet-scale augmented reality platform called the Niantic Real World Platform. This includes a scalable engine for shared state and user interactions that, with Pokemon Go, has already supported hundreds of millions of users. It is supported by a client platform that includes mapping, security, and AR capabilities.

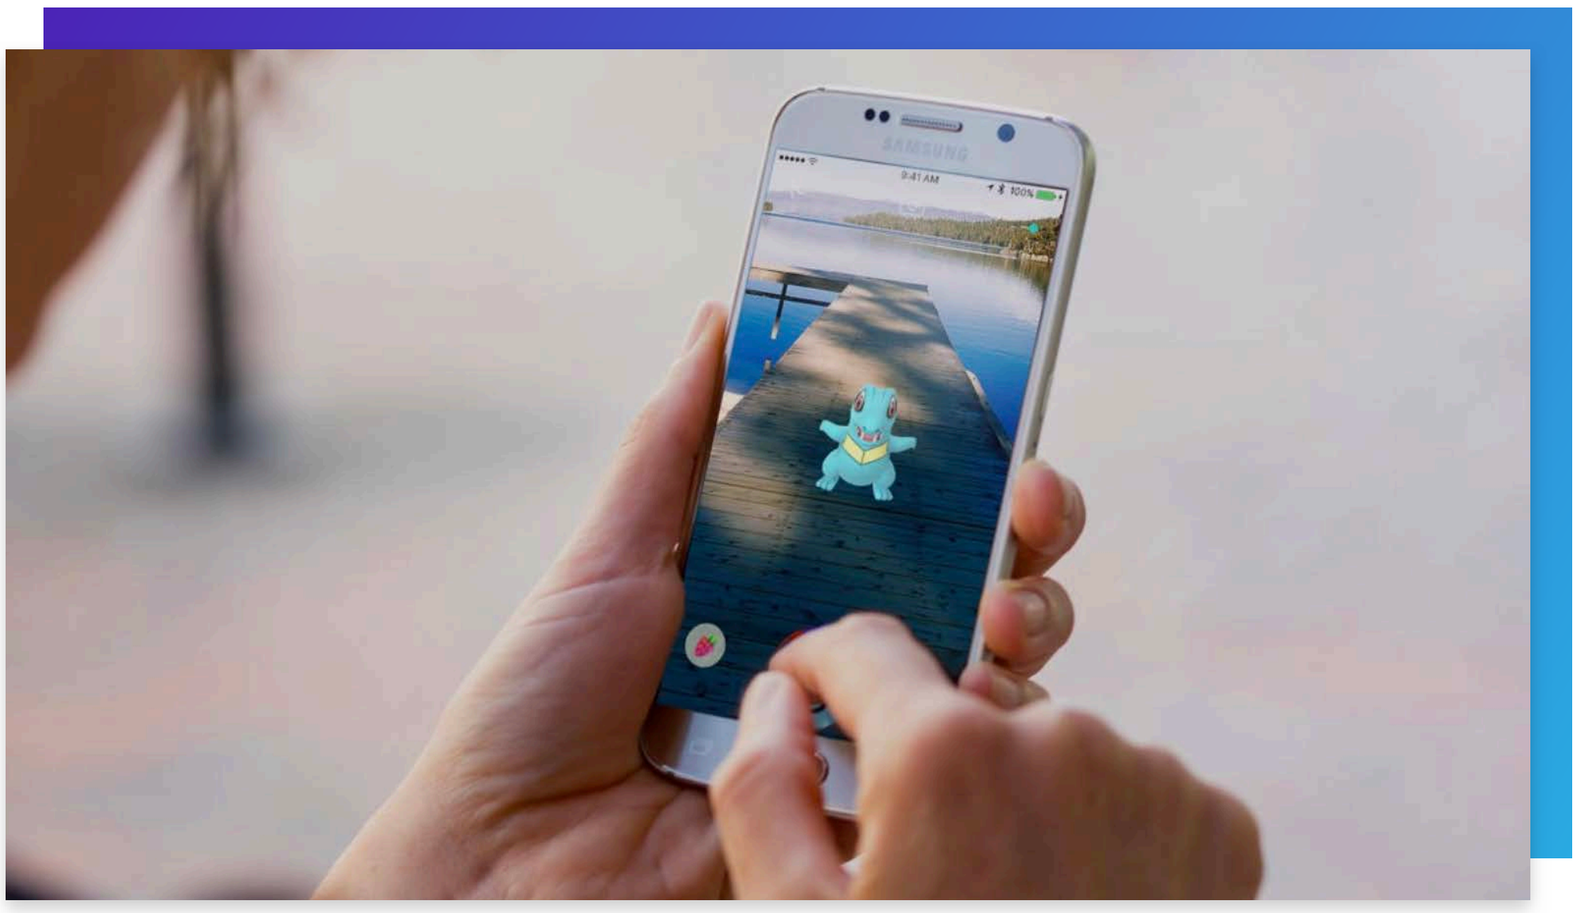

Image source: https://www.nianticlabs.com/img/index/screenshot-pokemongo2.jpg

2 2.2 Indoor Navigation Immersal

Use case for AR Cloud based augmented reality on mobile devices: Indoor navigation

Immersal launched an indoor navigation app in January 2017 for Messukeskus, the Helsinki exhibition center in Finland. It was based on ArUco markers (100 physical two-sided markers hanging on the ceilings in two large halls). Since using many physical markers is not a very scalable solution, Immersal developed markerless tracking based on computer vision and combined it with ARKit and ARCore. Immersal is now able to position virtual content to the real world within 1-3 cm (0.4-1.2 inch) accuracy. Immersal AR Cloud SDK will be released in early 2019.

Immersal is working with two international retail clients in Finland to provide them an AR Cloud based indoor navigation system (among other features). By scanning (i.e., constructing a point cloud) the (large) spaces beforehand, dynamic navigation is provided to points of interest (POIs)–the user can be anywhere on the premises and start navigating, and the navigation mesh will update accordingly and show the shortest path to a POI. The navigation arrow is responsive to occlusion, thus improving the UX. It also works across floors, which is important in a multi-storey venue. The indoor navigation is used to find departments, restrooms, restaurants & cafés, information desks etc. The 3D scanning of a venue can be done with a smartphone, and in several phases if needed.

Product finding is essentially the same feature as indoor navigation. If the retail store knows exactly where all their products reside in, we can tap into their systems and make finding products a breeze by using AR. Also available is product highlighting; e.g. show only vegan products in a supermarket and create a virtual path for the user to collect them faster.

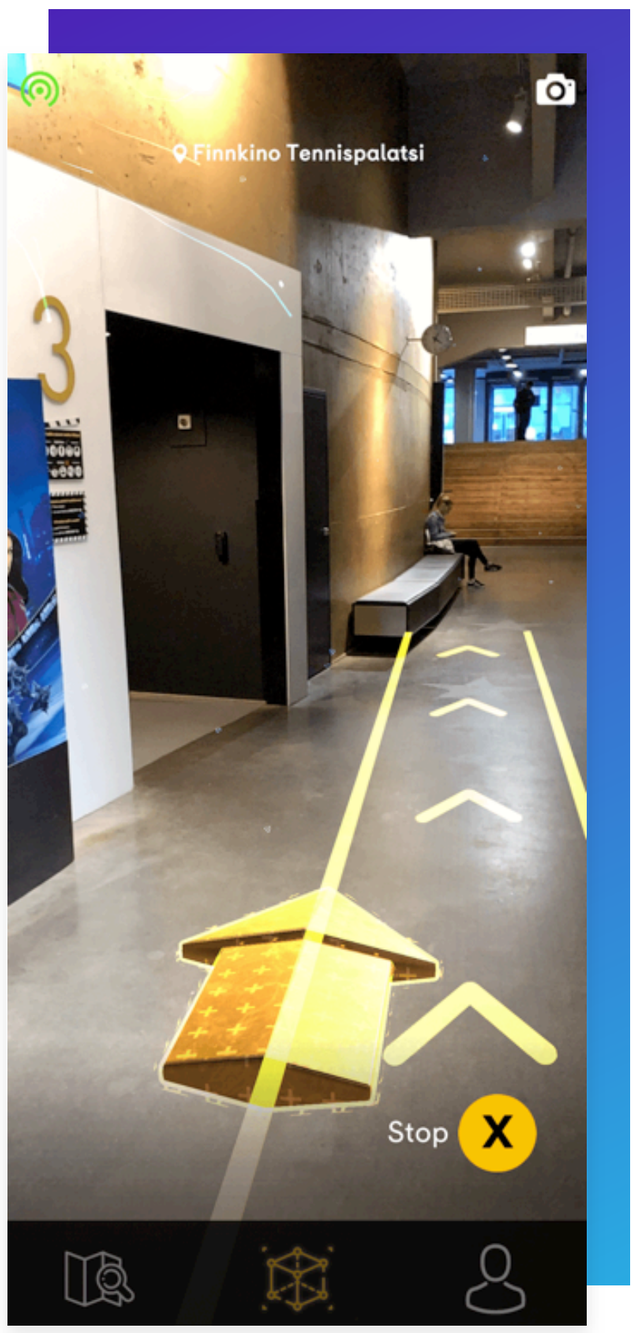

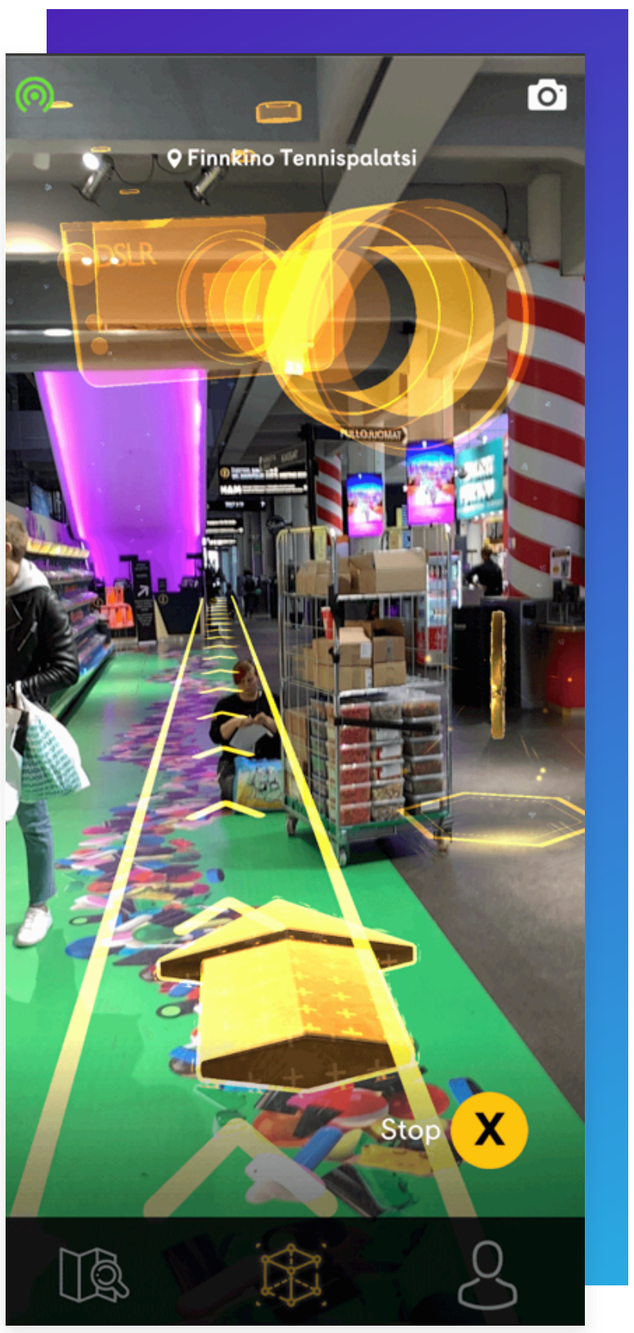

2 2.3 Productivity: YOU^{AR}

YOU^{AR} is building a dedicated AR Cloud platform, replete with tools that discover, share, create, and remix AR. All content on YOU^{AR} is accessible regardless of the device, whether it's running on iOS, Android, or an HMD. Powered by a planet-scale, multi-map CV and asset delivery framework, YOU^{AR} allows users to capture the "last meter" of the AR Cloud. Satellites, LIDAR drones, and GPS can only provide so much detail — but the mobile device, *already in everybody's pockets*, enables the capture of on-demand crowd-sourced spatial data with sub-meter accuracy.

"YOU^{AR}... addresses some of the most vexing problems in AR, including convergent cross-platform computer vision (real-time interaction between ARKit and ARCore devices), interactivity of multiple AR apps in the same location across devices, real-time scene mapping, geometric occlusion of digital objects, localization of devices beyond GPS (the AR Cloud), and the bundle drop of digital assets into remote locations. Together, this represents a heretofore unheard of stack of AR and computer vision features we have yet to see in AR, and could revolutionize the development of new apps."

— Charlie Fink, *Forbes*

Key AR Cloud Elements:

- Re-localization of devices within real space (providing 6 degrees of freedom)
- Integration of disparate point clouds into a comprehensive "multi-map" feature set
- Accurate and persistent positioning of cross-device dynamic digital content.
- Massive, synchronized, multi-user experiences between local and remote AR Cloud users
- Open system supporting solutions for geometry acquisition for real-time occlusion and collision of digital content within real spaces
- Social sharing of AR content

In 2019, YOU^{AR} has partnered with industry experts to roll out a version of its social AR Cloud platform for the manufacturing industry. This first step improves AR accessibility, providing training and real-time IoT device visualization to factory floor. Partner companies will create, use, and gamify training modules; they will drop location-specific AR notes to help employees or managers collaborate; they will let workers visualize IoT device data where they need to in an ordinary work day.

Being able to harness the power of AR, without constraining the choice of hardware involved, provides an immediate impact to a manufacturer's bottom line. Excited by this initial rollout, YOU^{AR} plans to develop its platform with these select partners before releasing a consumer version of its AR Cloud platform at a later date.

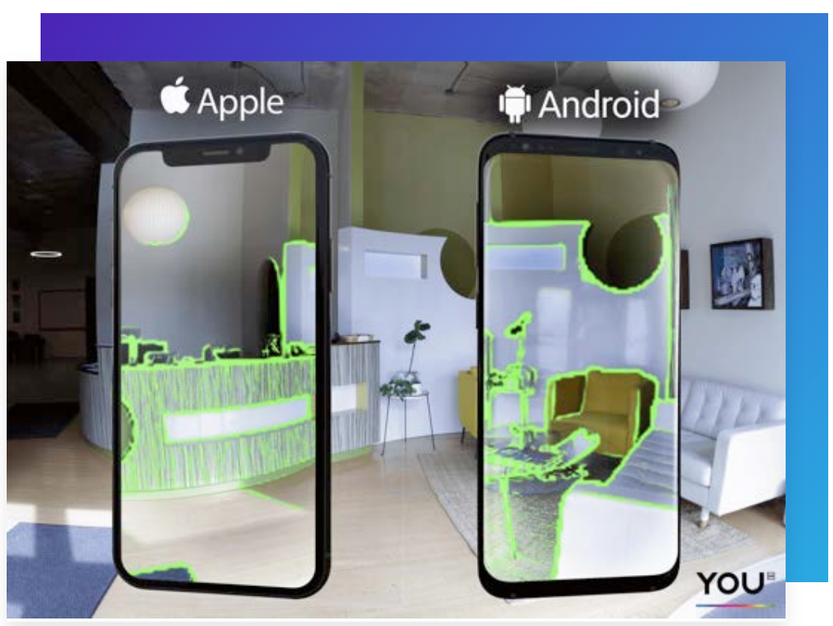

2 2.4 Social and Gameplay: Ubiquity6

Building compelling, interactive, persistent AR gameplay has the potential to be the first AR use case that uses the AR Cloud and drives large consumer adoption.

In August 2018, Ubiquity6 partnered with the SFMOMA to put on an interactive AR exhibit in the museum through PlaySFMOMA, the museum's initiative that encourages the development of avant-garde and experimental games. Over 100 people were able to simultaneously localize together and view AR art (in the style of René Magritte) as well as play a voxel building game together. This was the largest number of people that had ever viewed AR content together -- all localized and interactive in the same space with these persistent assets.

The AR Cloud will become even more immersive for gameplay as the technology advances. Games that involve an even larger number of users or more complicated gameplay mechanics will push this use case even further. The interactions could be shared between the users localized in the AR content, as well as others viewing or playing remotely. Like other technologies before it, the success of getting a massive following of users playing games that utilize the AR Cloud can also greatly push forward the other use cases that we expect to come soon after.

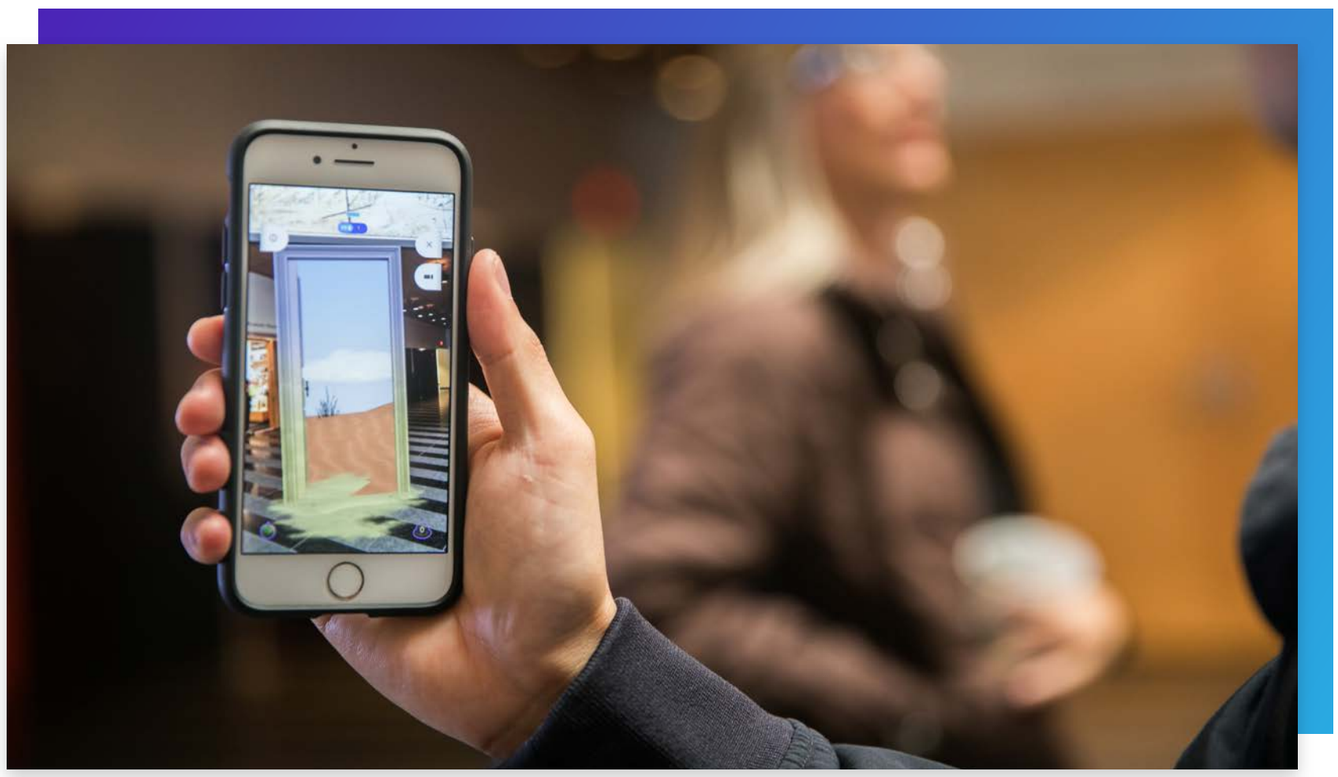

Image source: Ubiquity6

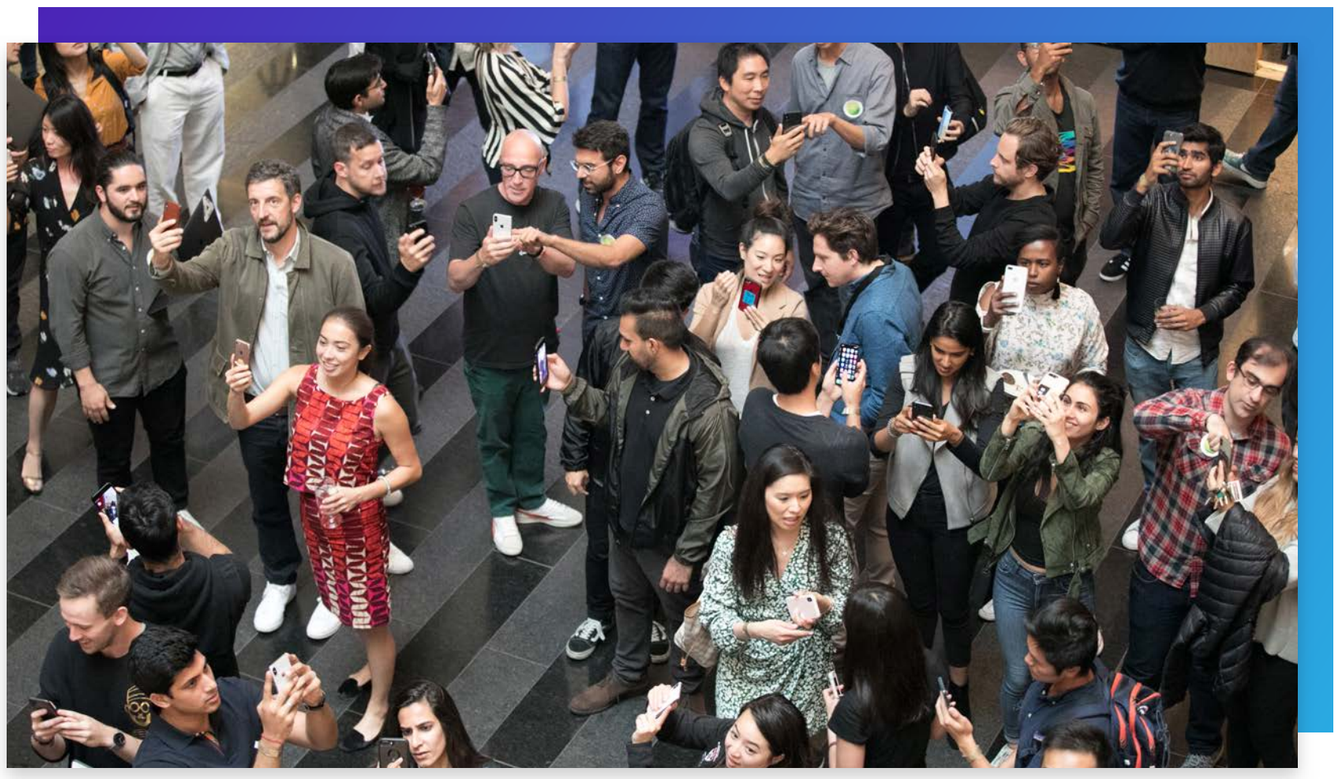

Image source: Ubiquity6

2 2.5 Events: Geogram

Spatialized event modeling platform that turns locations into smart spaces using the AR Cloud

A 3D editor, CMS and publishing desk, Geogram enables event managers to remotely design, schedule and manage physical and virtual environments, objects and events at real world locations. The result is a "smart space" that can be continuously updated: real events can be scheduled at real spots at the location, such as tents or kiosks, to which users can be navigated using AR wayfinding; virtual content can be placed in the smart space and triggered by user actions such as proximity or gesture. Geogram has partnered with Now Communications to roll out a smart space at Dundas Square - the most popular location for outdoor events in Toronto - in the Spring of 2019. Through the AR Cloud, users at Dundas Square can experience persistent multi-user AR with shared gameplay, navigation and other interactions with virtual objects.

These same features are bundled into developer tools that can be used to make any indoor or outdoor location a smart space for events such as conferences, trade shows, exhibits, concerts and festivals.

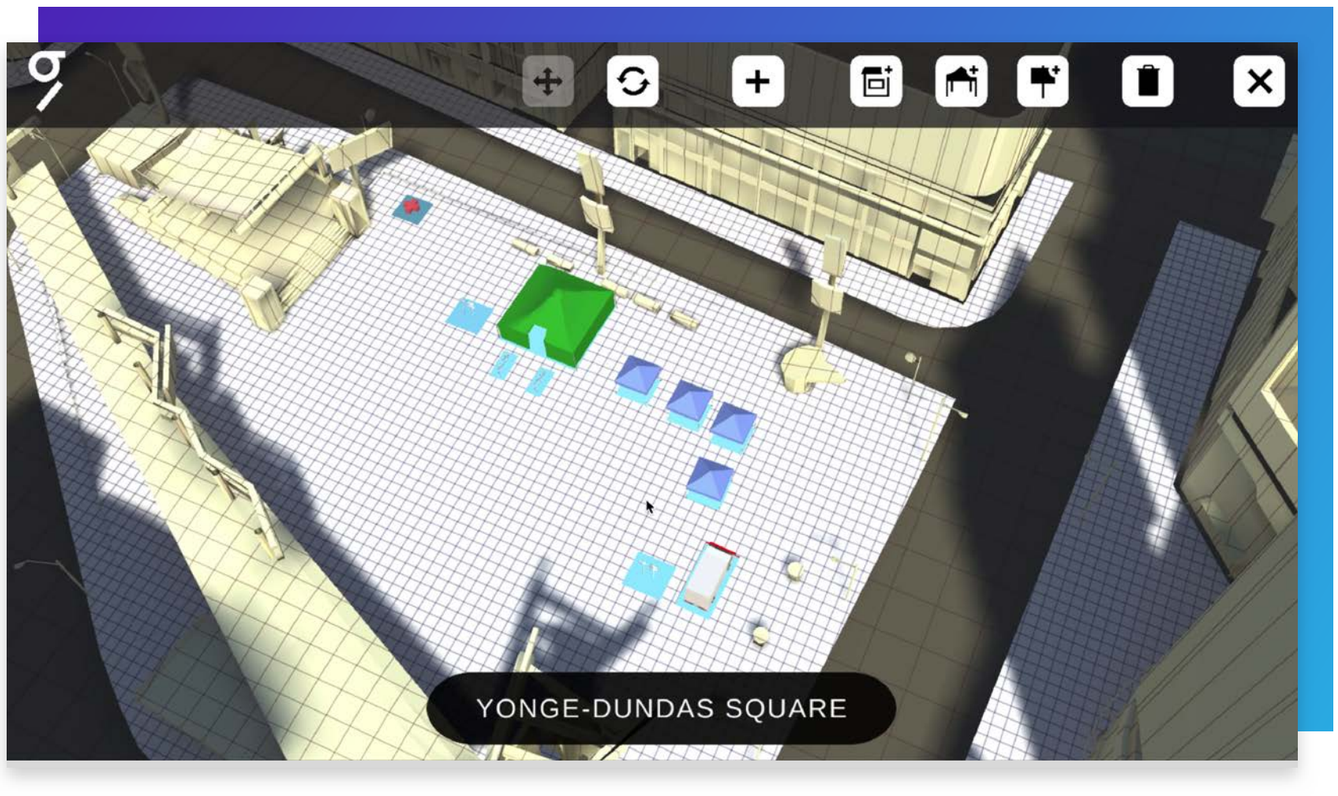

Image source: Geogram

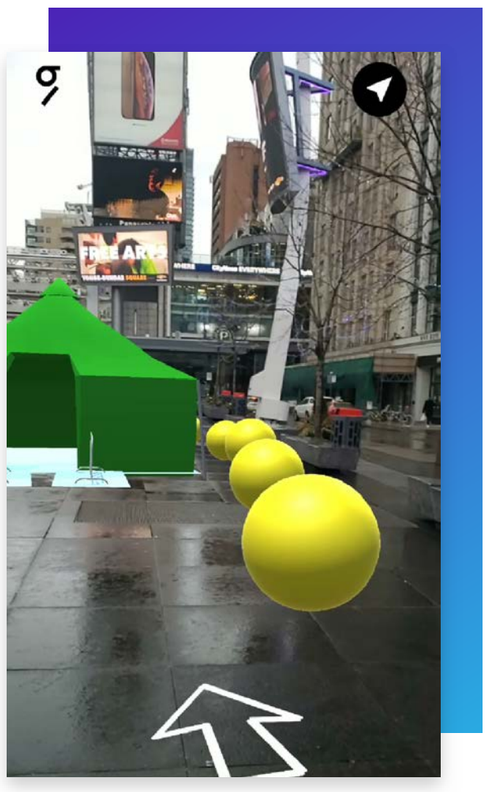

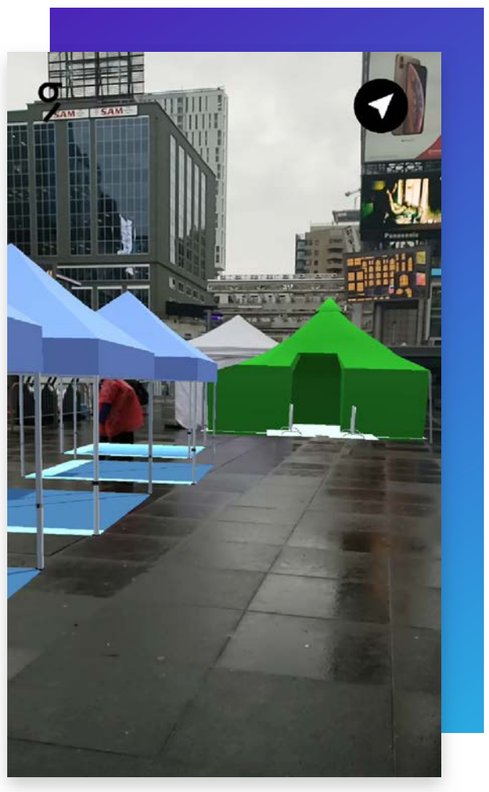

2 2.6 Location and Tracking: Fantasmo

Fantasmo is an 3D mapping mapping startup that has been building the Camera Positioning Standard (CPS) to give self-driving cars, robots and AR apps the ability to dynamically understand the real world.

Since mid-2018 Fantasmo has been working on the use case of scooter accountability - where the CPS is used to more precisely track the location of scooters for tracking pick up and drop off than the less accurate GPS. They also work with private enterprise customers that operate body-worn devices and robots in private locations. Jameson Detweiler, CEO and Co-Founder of Fantasmo explains that the vision is to build "a decentralized, cross-platform real time 3D map of the world that reflects the legal and social contracts we have around the spaces we occupy."

Companies are taking notice. In TechCrunch, Lime's VP of Global Expansion Caen Contee was quoted as saying, "We believe that mapping, 3D-imaging, and AR are amazing tech levers to drive scooter behavior change in a scalable, low-infrastructure way. We have been investing in this area among others in our quest to create the best urban mobility experience."

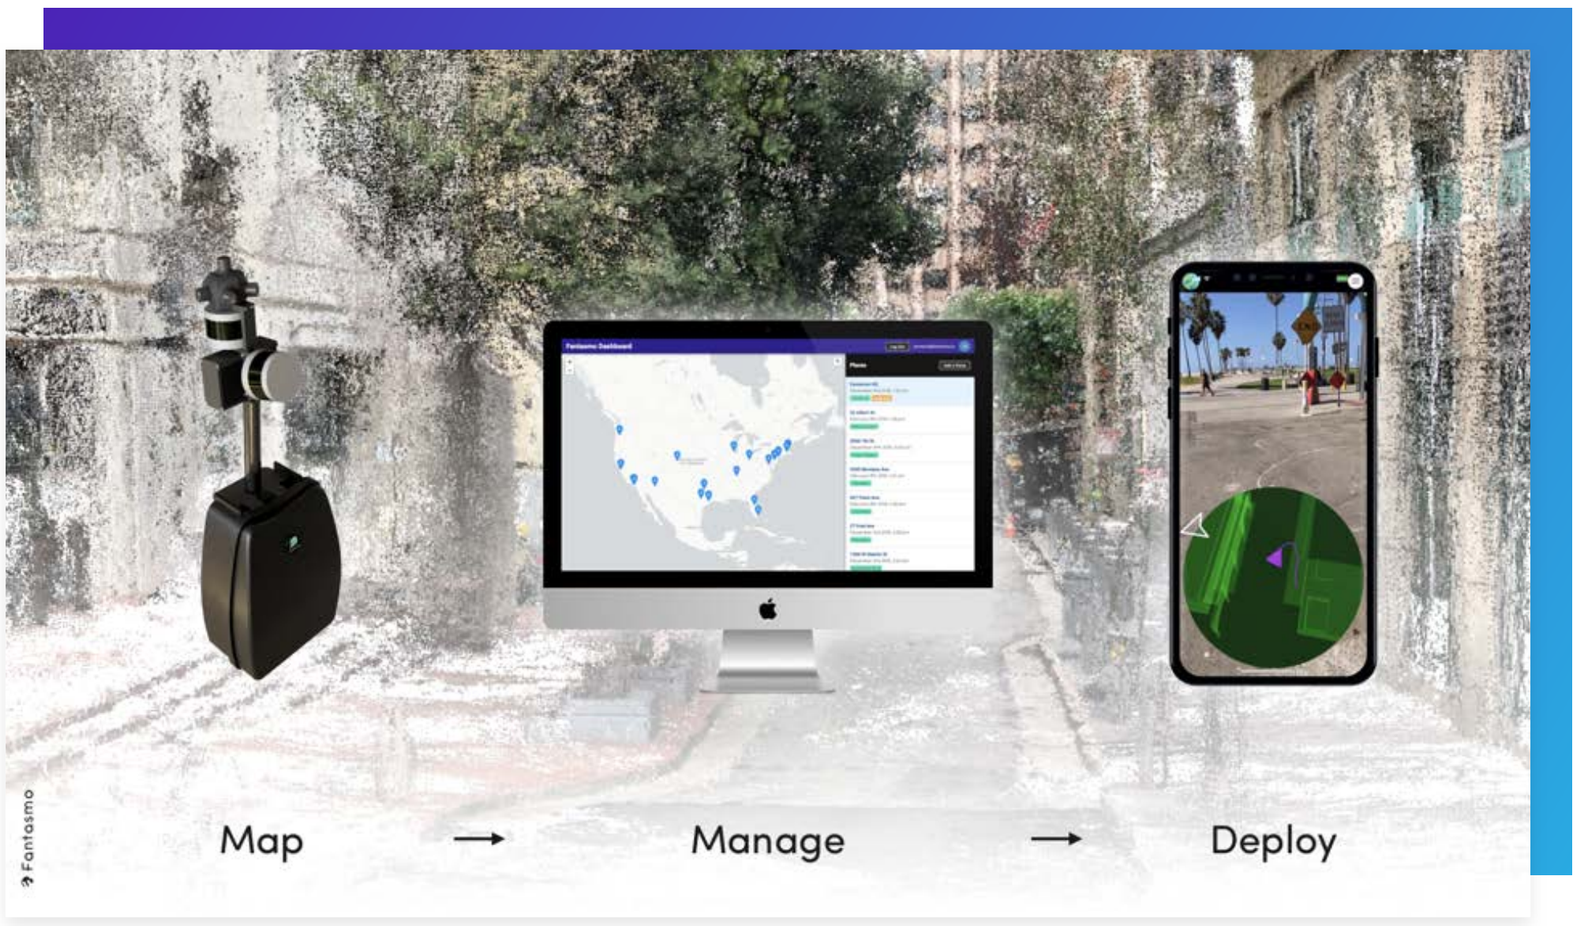

Image source: Fantasmo/NextReality

2 2.7 AR real estate: SuperWorld (superworldapp.com)

SuperWorld is a social Augmented Reality real estate platform on the blockchain using AR Cloud.

How does it work?

Utilizing the ERC721 standard, a user can go on to the SuperWorld platform and purchase the rights to any available plot of land on the real world map interface. When a user buys a plot of land they are buying an ERC721 token that corresponds to the unique latitude and longitude polygon they have selected. For example, if someone were to buy a section of Times Square in NYC, that section of land has an intrinsic string of code embedded into it and can only be associated with one person or to be more specific, one account address at a time. The owner of this AR real estate can do whatever they chose with it. Whether that be holding it long term and sharing in future ad revenue generated from AR ads placed on their property, or listing it for resale immediately — when someone buys a piece of AR real estate they are buying a unique, non-fungible piece of the blockchain.

Users can use the SuperWorld app on iPhone and Android phones to place AR objects, photos, videos, messages and play games using AR cloud. SuperWorld's mobile app is the first integration with the SuperWorld AR real estate ecosystem. We built it to live test our AR features. Eventually, the app itself will also evolve into an SDK for other developers to build upon and integrate into our AR Real Estate marketplace.

The team plans on creating an AR ad network upon which brands and advertisers will be able to access and serve ads on real world locations through any AR app connected to the SuperWorld platform. Similar to Google Ad Sense, brands will be able to create and run entire ad campaigns but instead of their ads appearing on a selection of high traffic websites, they will appear in high traffic locations around the world using the AR Cloud.

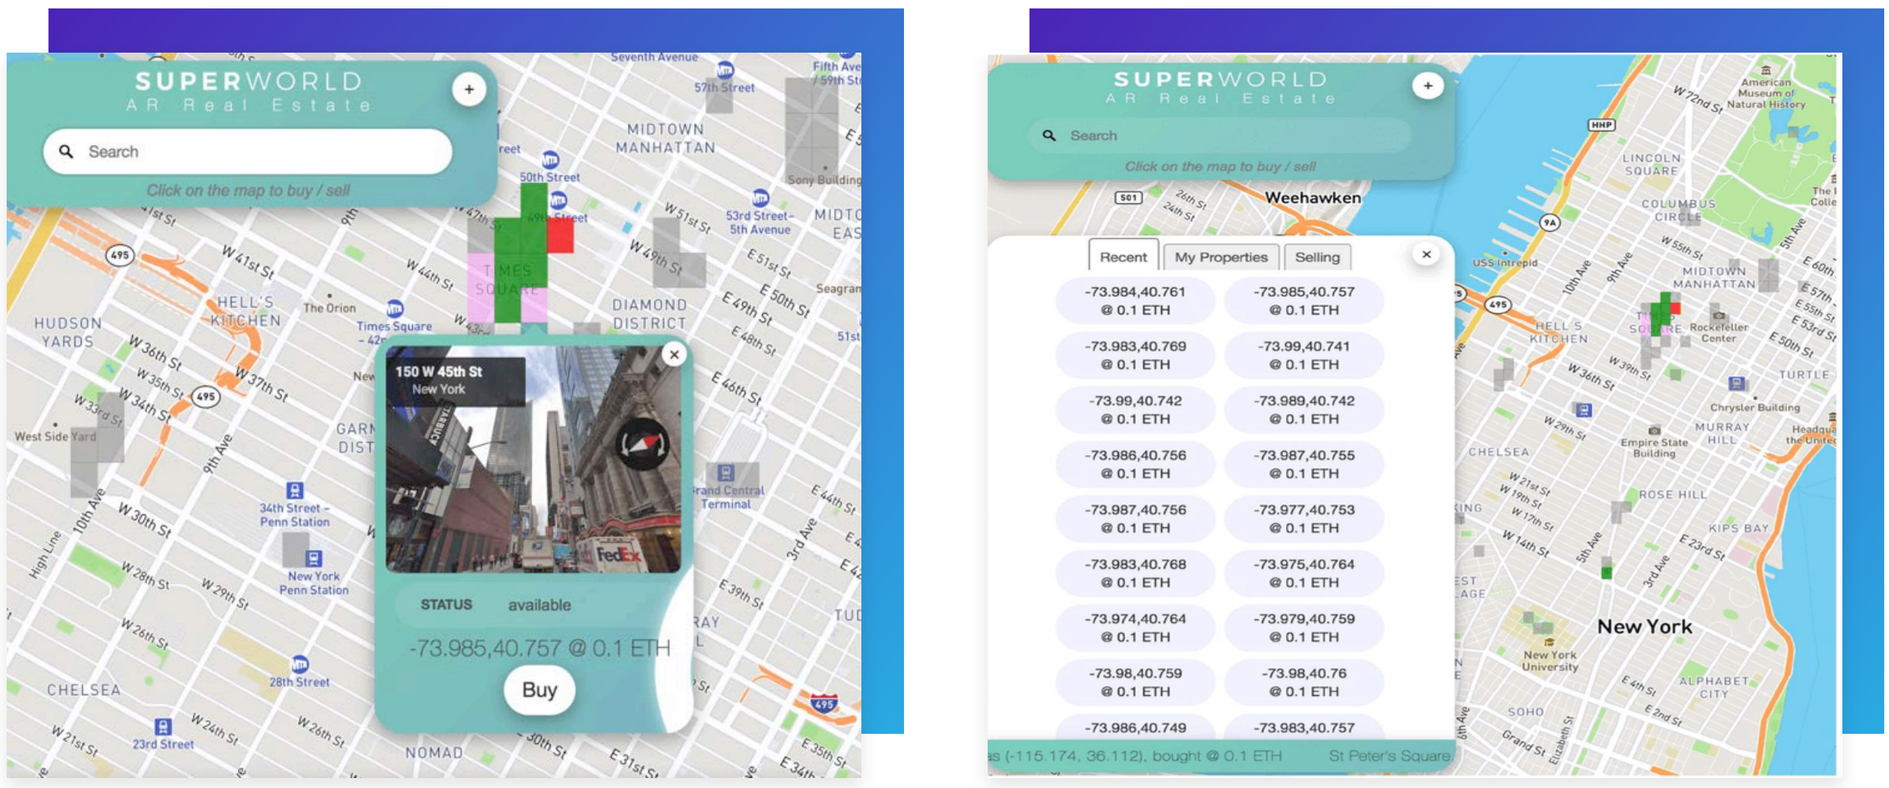

Image Source: SuperWorldapp.com

3 Conclusion

The AR Cloud will be one of the most important infrastructures in the history of computing–as not only a spatial map of the world, but a living history or soft copy of the world.

While the AR Cloud promises an exciting future, it will be incredibly challenging to build. Creating a real-time infrastructure requires robust computer vision technology that continuously collects and updates the spatial data while also differentiating between static and dynamic elements. Further, for the AR Cloud to work effectively and efficiently, a network of smartphones, smart glasses, CCTV cameras, self-driving cars, drones and satellites all need to agree on how data is captured, processed, and interpreted. That said, like the world wide web, once built, its application will extend well beyond Augmented Reality into drone routing, self-driving vehicles, other autonomous robots and more.

Participate in the VR/AR Association AR Cloud Committee (http://www.thevrara.com/ar-cloud) by visiting www.thevrara.com

About VRARA

VRARA is the VR/AR Association, an international organization designed to foster collaboration between innovative companies and people in the VR and AR ecosystem that accelerates growth, research and education, and develops best practices and guidelines. VRARA has over 4200 companies and 25,000 professionals registered, over 50 chapters globally, and 20 industry committees. VRARA programs & initiatives are designed to accelerate anyone's growth, knowledge, and connections.