

# How Facebook researchers' realistic simulations help advance AI and AR

Replica and AI Habitat are key to developing experiences like social presence and intelligent assistants

June 14, 2019



As the AI bot moves through the room, it navigates past two couches, a coffee table, and some small tan chairs. Zipping over the area rug, it heads toward the keys that researchers asked it to find.



In this case, though, the environment is a digital simulation that is part of Replica, a research project created by Facebook Reality Labs (FRL). Replica is a photo-realistic re-creation of 18 sample spaces, such as an office conference room and a two-story home set up by researchers.

## Recommended Reading

[Facebook is building the future of connection with lifelike avatars](#)

---

[How Facebook's research lab is thinking about AI](#)


---

[Speechless? Here's how AI learns to finish your sentences](#)

---

FRL made these virtual spaces to help AI researchers teach machines about the real, physical world — an important step in developing more capable real-world assistants as well as next-gen augmented reality (AR) and virtual reality (VR) experiences. The idea is that if researchers can train an AI system to locate a virtual missing set of keys in a lifelike digital living room, it will eventually enable a real-life robot to do the same with real keys in a real room, too. And if AR and VR applications can learn how to interact with different physical environments, then one day we'll be able to use our [photo-realistic digital avatars](#) to drop in on a family birthday party halfway across the world.

But researchers believe these systems will learn best if the simulated environments capture the subtle details — like mirror reflections and rug textures — necessary to make them virtually identical to the real thing. That's why FRL created Replica.





This video shows a Replica simulation alongside the actual real-world space. (The image on the right is the digital re-creation.)

“The Replica data set sets a new standard in the realism and quality of 3-D reconstructions of real spaces,” says Julian Straub, an FRL Research Scientist who worked on creating Replica. Straub studied electrical engineering in Germany before earning his PhD in computer science at MIT and eventually joining FRL to work on machine perception. As described by its Chief Scientist, Michael Abrash, [FRL’s mission](#) is to develop the technologies needed to establish AR and VR as the next computing platform. Projects like Replica will play an important part in realizing that vision.






This flyover shows a digital Replica environment created by Facebook Reality Labs.


The accuracy and fidelity of Replica's digital re-creations come from the combination of a well-engineered and integrated camera rig with a high-accuracy depth capture system, a state-of-the-art simultaneous localization and mapping (SLAM) system, and a dense reconstruction system. FRL's high-accuracy depth capture system, which uses dots projected into the



scene in infrared, serves to capture the exact shape of big objects like tables and chairs and also smaller ones, like the remote control on the couch.

The custom-built SLAM and dense reconstruction system transform the raw video streams captured from the camera rig into Replica re-creations of the real spaces that even a careful observer might think are real. (More details can be found in the Replica [data set white paper](#) as well as in [the team's 2018 SIGGRAPH conference paper describing the mirror and glass reconstruction system.](#))





A virtual robot (known as an embodied agent in the research community) is spawned in an unfamiliar Replica environment. It starts at a random position, shown here as a blue dot, and is then asked to navigate to a destination (the red dot). The agent is given instructions relative to its starting position — e.g., "Go 10m north and 15m west" — but it has no map to guide it. The agent must use only its sensory input (a regular RGB camera and in some cases also a depth camera) to accomplish its goal.

## Practicing a task millions of times in just one hour

Replica can be loaded up in AI Habitat, a new open platform for embodied AI research. Facebook AI created AI Habitat to be the most powerful and flexible way for researchers to train and test AI bots in simulated living and working spaces. AI Habitat allows researchers to put a bot into a Replica environment, so it can learn to tackle different tasks, like "go check if my laptop is on my desk in the kitchen." These chores are simple for humans, but for machines to master them, they must recognize objects, understand language, and navigate effectively. Today's machines — like robotic vacuums, for example — can respond to commands, but they don't understand and adapt to the world around them as people do. AI Habitat will help researchers develop bots that understand the physical world. But it is also an important research tool for creating next-gen AR experiences that begin to merge the physical and digital worlds. If we can teach an AI system to understand the physical space around you, we might one day be able to use it in combination with AR glasses. For example, it could help us to place your grandma's digital avatar in the seat next to you or to display digital reviews right next to a restaurant or store as you walk by.


Replica provides realistic 3D data, and AI Habitat provides simulation with speed and flexibility. While other simulation engines commonly run at 50 to 100 frames per second, AI Habitat runs at over 10,000 frames per second (multi-processor on a single GPU). This enables researchers to test their bots much more quickly and effectively — an experiment that would



take months on another simulator would take a few hours on Habitat. Facebook AI research intern Erik Wijmans, who is also a PhD student at Georgia Tech, and AI Resident Bhavana Jain used the system to do state-of-the-art research, training their bot with over a billion frames of experience. Using a rough estimate of how quickly people can look around and move in the real world, that would be the equivalent of more than 30 years of experience. A virtual bot can also bump into countless walls and make other mistakes as it learns, without any risk of doing real-world damage.

Facebook is now open-sourcing AI Habitat and releasing its Replica data set, so anyone in the community can build on it, try new approaches, compare their results, and learn from others' work. (Technical details on Habitat are [available here](#), and the Replica environments can be downloaded [here](#).) This kind of open sharing of information between researchers at different companies and organizations has been key to recent advances in AI technologies like [natural language understanding](#), [computer vision](#), and [embodiedQA](#), and researchers at Facebook AI and FRL believe the same will be true here.

To establish performance benchmarks that can be used by everyone in the field, Facebook AI also recently created the Habitat Challenge. The contest invited engineers and researchers from across the AI community to find the best way for bots to complete a particular navigation task in AI Habitat. "AI Habitat offers close to real-world experience for learning navigation," says one of the challenge participants, Dmytro Bobrenko.





The Replica data set identifies and labels the objects found in its virtual spaces, assigning a different color to different categories of objects, like “chair” or “wall.” AI researchers can use this “semantic segmentation” data to help develop smarter systems.

Together, these technologies will one day enable machines to learn to operate intelligently in the real world rather than just on our smartphones or laptops, says Dhruv Batra, a Facebook AI Research Scientist and Georgia Tech professor who leads the Habitat team. He and his colleagues call this the shift from “internet AI” to “embodied AI.” It means teaching machines using not just static data sets (like photos of cars) but also interactive environments (such as a simulated parking lot full of virtual cars that an AI bot can explore and examine). Batra and many other AI researchers believe this kind of interaction is necessary for building a new wave of smart tools to help us in the physical world as well as in the digital one.

## Building the tools to create ‘social presence’






By training systems through AI Habitat’s advanced, open-platform simulations, researchers can make progress on embodied AI technologies that, until now, have largely remained in the realm of science fiction. Batra foresees not just intelligent assistants but also tools to help the visually impaired better navigate their surroundings, for example.

For Richard Newcombe, a Research Director at FRL, one of the most exciting future applications is bringing “social presence” to the physical world around us. VR today enables people to share a virtual space with a friend who is hundreds of miles away. Newcombe is working to bring a new level of realism to the experience and to make social presence possible in everyday life, too, through the use of AR glasses. With this technology, you’ll one day be able to see and interact with that friend just as if he or she were sitting on the couch next to you. To create this kind of social presence, AI systems need to know how to make the digital version of your friend interact naturally and realistically with the actual physical couch and room around you, or put you in a simulated environment that looks and seems real.

“Much as the FRL research work on [virtual humans](#) captures and enables transmission of the human presence, our reconstruction work captures what it is like to be in a place; at work, at home, or out and about in shops, museums, or coffee shops,” Newcombe says. He is passionate about developing technology that can sense and understand the state of the world. He started work in the field with an apprenticeship at age 16 and went on to study robotics, computer vision, and machine learning at the University of Essex before earning his PhD at Imperial College London. He joined Facebook four years ago to start the research and incubation team dedicated to the future of [machine perception](#) for AI and XR applications, and the launch of Replica represents an important step forward in making this a reality.





Another Replica environment captures details like the electrical outlets on the wall and objects behind the glass doors of a bookcase.

## A responsible, open approach

Building experiences like social presence will require additional breakthroughs in hardware as well as continued advances with training resources like Replica and AI Habitat. But there are also important privacy and security considerations, says Newcombe.



“We must be incredibly diligent in generating the reconstruction, scene understanding, and AI reasoning systems,” he says. Researchers and engineers as well as outside experts and the general public will have to collaborate to work through the social and personal consequences of a potentially transformative new technology. To do so, they need to know what is possible and to share updates on their progress so others can be part of that public discussion. At F8, Facebook’s keynote speakers discussed recent work on [ethical design](#) and [addressing bias](#), which will be important as research progresses on AR experiences and embodied AI.

With the Replica scans, the data was anonymized to remove any personal details (such as family photos) that could identify an individual. In building this 3D reconstruction technology, the FRL researchers also made sure they created a robust system for handling and storing data even before they started scanning. For example, the data is stored securely in a server that can be accessed only by a limited number of researchers, and the team has regular reviews with privacy, security, and systems experts to make sure they are following protocol and implementing the latest and most rigorous safeguards possible. The scans are made available to the broader research community only after these steps are completed.

It will take many more technological breakthroughs before experiences like AR social presence and advanced AI assistants are a reality. For example, Facebook AI researchers will explore ways to build realistic physics modeling into AI Habitat so an AI bot can learn what happens when it knocks a virtual glass off a virtual table. As this work progresses, the Replica and AI Habitat researchers believe these projects will play an important part in Facebook’s future. By enabling the next generation of embodied AI, these technologies will unlock the potential of AR glasses — giving people a better understanding of their world and helping them to connect, work, and collaborate in powerful new ways.

“Using AR glasses as the platform, social telepresence and useful AI assistants will help you be the most effective you ever, and most engaging in the world however you want to,” says Newcombe.

