

SPEAR: A Simulator for Photorealistic Embodied AI Research

Mike Roberts^{1,2}, Renhan Wang³, Rushikesh Zavar⁴, Rachith Dey-Prakash²,
Quentin Leboutet², Stephan R. Richter², Matthias Müller², German Ros⁵,
Rui Tang³, Stefan Leutenegger^{6,7}, Yannick Hold-Geoffroy¹, Kalyan Sunkavalli¹,
and Vladlen Koltun²

¹Adobe Research ²Intel Labs ³Manycore Tech Inc ⁴Adobe ⁵NVIDIA
⁶ETH Zurich ⁷Imperial College London

<https://github.com/spear-sim/spear>

Abstract. Interactive simulators have become powerful tools for training embodied agents and generating synthetic visual data, but existing photorealistic simulators suffer from limited generality, programmability, and rendering speed. We address these limitations by introducing *SPEAR: A Simulator for Photorealistic Embodied AI Research*. At its core, SPEAR is a Python library that can connect to, and programmatically control, any Unreal Engine (UE) application via a modular plugin architecture. SPEAR exposes over 14K unique UE functions to Python, representing an order-of-magnitude increase in programmable functionality over existing UE-based simulators. Additionally, a single SPEAR instance can render 1920×1080 photorealistic beauty images directly into a user’s NumPy array at 73 frames per second – an order of magnitude faster than existing UE plugins – while also providing ground truth image modalities that are not available in any existing UE-based simulator (e.g., a non-diffuse intrinsic image decomposition, material IDs, and physically based shading parameters). Finally, SPEAR introduces an expressive high-level programming model that enables users to specify complex graphs of UE work with arbitrary data dependencies among work items, and to execute these graphs deterministically within a single UE frame. We demonstrate the utility of SPEAR through a diverse collection of example applications: controlling multiple embodied agents with distinct action spaces (e.g., humans, cars, and robots) across several in-the-wild UE projects; rendering photorealistic city-scale environments; manipulating UE’s procedural content generation systems; rendering synchronized multi-view images of detailed human faces; coordinating an interactive co-simulation with the MuJoCo physics simulator; and editing scenes with natural language via an AI coding assistant.

1 Introduction

Interactive simulators have become a critical layer of scientific infrastructure, participating in major breakthroughs in reinforcement learning (e.g., world-champion performance in competitive e-sports games [56, 75, 76, 82] and drone racing [30–32, 43, 44]), sensorimotor control [60, 61, 88], autonomous driving [18,




Fig. 1: SPEAR is a Python library that can connect to, and programmatically control, any Unreal Engine (UE) application via a modular plugin architecture. SPEAR exposes over 14K unique UE functions, representing an order-of-magnitude increase in programmable functionality over existing simulators. We demonstrate the flexibility of SPEAR by using it to control 6 distinct embodied agents (each with a different action space) across several Epic Games sample projects: a person and a car from CITYSAMPLE (top); a flying robot from STACKOBOT (bottom far left); multiple agents in a resource collecting game called CROPOUTSAMPLE (bottom center left); as well as a person with parkour skills and a quadruped robot from GAMEANIMATIONSAMPLE (bottom right).

[53], dexterous manipulation [55], and quadruped locomotion [24, 39, 50]. Additionally, photorealistic synthetic datasets (e.g., [62]) are now being used to train and evaluate large-scale foundation models (e.g., [12, 20, 51, 69, 73, 85]), as well as state-of-the-art methods for 3D and 4D reconstruction [49, 77, 78], segmentation [35, 59], depth estimation [33], and controllable video synthesis [13].

Motivated by the parallel goals of interactivity and photorealism, the computer vision, robotics, and embodied AI communities are increasingly leveraging photorealistic simulators that are built on top of industrial-strength game engines like Unreal [8, 16, 18, 40, 47, 48, 58, 67, 86, 87], Unity [7, 21, 28, 36], and Panda3D [5, 80], as well as closed-source commercial simulation platforms like IsaacSim [3, 22, 23, 42, 81]. Among these alternatives, the Unreal Engine (UE) [8] is a compelling platform due to its state-of-the-art real-time rendering [29, 54, 79] and procedural content generation systems [38], its proven track record in commercial game development [10, 11], and because it is completely open-source.

However, existing simulators that build on UE have several important limitations. First, all existing UE-based simulators expose relatively limited hand-crafted Python interfaces (i.e., consisting of hundreds of functions). Second, existing UE-based simulators can incur significant communication overhead when returning large blocks of data (e.g., high-resolution images) to a user’s Python program. For example, in the UE-based simulators we tested, we found that returning a 1920×1080 image to Python can be $20\text{--}35\times$ slower than rendering the same image to the viewport in a standalone UE application with no Python communication. Third, most UE-based simulators are distributed as large monolithic applications, in some cases requiring their own custom fork of the 10+ million




Fig. 2: SPEAR includes a customizable camera sensor that can render 1920×1080 photorealistic beauty images (left) directly into a user’s NumPy array at 73 frames per second – an order of magnitude faster than existing UE plugins – while also providing ground truth image modalities that are not available in any existing UE-based simulator. For example, the SPEAR camera sensor can render all of the image modalities in the Hypersim dataset [62], i.e., depths, surface normals, instance and semantic IDs (right top), and a non-diffuse intrinsic image decomposition (right bottom), as well as material IDs and physically based shading parameters (see our supp. material).

line UE codebase, rather than as modular plugins. As a result, it is not straightforward to integrate these simulators into existing projects, and likewise, it is not straightforward to integrate third-party assets into existing simulators.

In this work, we introduce SPEAR, a simulator for **photorealistic embodied AI** research that addresses all of the limitations described above (see Figures 1 and 2). Throughout the design and implementation of SPEAR, our key technical insight is to expose a comprehensive C++ interface for interacting directly with UE’s runtime reflection system [9, 45], thereby enabling our Python library to find classes, call functions, and manipulate variables on objects dynamically at runtime using strings as keys (i.e., without requiring hand-crafted wrapper code for each class, function, and variable). By hooking directly into UE’s reflection system, SPEAR exposes a significant amount of underlying UE functionality (over 14K functions) while maintaining a modest code footprint (see Table 1).

On top of our low-level reflection interface, we introduce an expressive high-level programming model that enables users to call any UE function as though it was a native Python function, provided the UE function is visible to the reflection system (see Figure 3). For improved efficiency, we provide an asynchronous variant for every function in SPEAR that avoids synchronizing with UE, and we implement a mechanism for passing NumPy arrays back and forth between UE functions and user Python code without requiring any data copying.

Together, these novel technical contributions enable an order-of-magnitude increase in programmable functionality over existing UE-based simulators, while also enabling us to render 1920×1080 photorealistic beauty images directly into a user’s NumPy array at 73 frames per second – an order of magnitude faster than existing UE plugins (see Table 2).

We demonstrate the utility of SPEAR through a diverse collection of example applications: controlling multiple embodied agents with distinct action spaces (e.g., humans, cars, and robots) across several in-the-wild UE projects; rendering photorealistic city-scale environments; manipulating UE’s procedural content generation systems; rendering synchronized multi-view images of detailed human faces; coordinating an interactive co-simulation with the MuJoCo physics simulator [72]; and editing scenes with natural language via an AI coding assistant.




Fig. 3: (a, b): We demonstrate several fundamental concepts in the SPEAR programming model with a simple example program that spawns a set of coordinate axes in an indoor environment. In our programming model, graphs of UE work are specified as transactions. In particular, the user specifies a transaction by defining a `begin_frame` context followed by an `end_frame` context. Within each context, the user specifies a graph of UE work simply by implementing it as Python code. Any C++ function or variable that is visible to UE’s reflection system (e.g., `SetActorScale3D`, `RootComponent`) can be accessed as though it was a native Python function or attribute. For improved efficiency, we provide an asynchronous variant for each function in SPEAR (e.g., `async.GetComponentLocation`) that avoids synchronizing with UE. (c): It is trivial to expose new C++ functions and variables to UE’s reflection system, and therefore to SPEAR, simply by adding a `UFUNCTION` or `UPROPERTY` annotation next to the function or variable in any C++ header file.

2 Related Work

Simulators for Embodied AI We discuss interactive simulators based on game engines and commercial simulation platforms in Section 1. Additionally, the embodied AI community has developed simulators for navigating through static scans of real-world environments [64, 65, 83], and manipulating fully interactive environments [41, 57, 68, 71, 84], using custom simulation stacks that typically prioritize simulation speed over photorealism. In contrast, we choose to build on top of UE to leverage its state-of-the-art rendering and procedural content generation systems, as well as its vibrant ecosystem of high-fidelity content.

Using Games as AI Simulators Spanning across the history of the medium, commercial video games have been successfully repurposed into AI simulators, from the earliest Atari 2600 games [15, 52], to early first-person 3D games [14, 25, 34, 88], to modern strategy [56, 75, 76], racing [82], and open-world exploration [19, 27, 37, 46, 60, 61] games. However, these simulators either require a hand-crafted scripting interface to be implemented by the game’s developers (e.g., [76, 82]); or are implemented as unintrusive wrappers that intercept graphics driver communication (e.g., [37, 60, 61]) and programmatically inject device input (e.g., [46]), and therefore offer limited opportunities to observe and manipulate the underlying game state. In contrast, we choose to implement our simulator as a set of modular plugins that provide comprehensive programmatic access to UE, and can be enabled in any UE application by adding a single-line declaration to a project file.

AirSim Single step	UnrealCV+ Batched commands	Habitat 2.0 Double-buffered observations	CARLA Sync and async modes	External physics simulator User-defined sub-stepping
<pre>def step(action): with begin_frame(): unpause() apply_action(action) with end_frame(): pause() with begin_frame(): obs = get_observation() with end_frame(): pass return obs</pre>	<pre>def step(actions): with begin_frame(): apply_action(actions[0]) apply_action(actions[1]) with end_frame(): obs = get_observation() pass return obs</pre>	<pre>def initialize(): with begin_frame(): pass with end_frame(): obs = get_observation() return obs def step(action): with begin_frame(): unpause() with end_frame(): future = asyncio.get_event_loop().create_future() apply_action(action) future.set_result(obs) return future.get()</pre>	<pre>def step_sync(action): with begin_frame(): apply_action(action) with end_frame(single_step=True): obs = get_observation() return obs def step_async(action): with begin_frame(): apply_action(action) with end_frame(): pass with begin_frame(): obs = get_observation() with end_frame(): pass return obs</pre>	<pre>def initialize(): with begin_frame(): disable_ue_physics() with end_frame(): pass def step(action): for i in range(num_substeps): my_phys_sim.step(action) with begin_frame(): for o in get_ue_objects(): if my_phys_sim.has_object(o.name): p = my_phys_sim.get_pose(o.name) set_ue_pose(o, p) with end_frame(): obs = get_ue_observation() return obs</pre>

Fig. 4: We demonstrate the flexibility of our programming model by using it to implement a variety of synchronization strategies found in existing simulators. For each strategy, we implement a simplified OpenAI Gym `step` function [17], which takes an `action` as input, steps a simulation forward, and returns an `observation` as output. In our programming model, it is straightforward to implement AirSim’s approach for single-stepping [67], UnrealCV+’s approach for batched commands [58, 87], Habitat 2.0’s approach for double-buffered observations [71], CARLA’s synchronous and asynchronous stepping modes [18], and co-simulation via an external physics simulator (e.g., [72]) with user-defined sub-stepping.

Scripting Interfaces to the Unreal Engine UE includes a powerful node-based visual programming environment known as *Blueprints* [1], which exposes nearly every class, function, and property that is visible to the reflection system, and can be used to script standalone applications. But Blueprints are unfamiliar to most AI practitioners, and they are stored in a proprietary binary data format that is incompatible with text-based developer workflows (e.g., diffing and merging tools, AI coding assistants). On the other hand, the Unreal Editor includes a self-contained Python environment that exposes a comparable amount of functionality to Blueprints [6]. However, this Python environment is intended for automating offline content-creation tasks inside the Unreal Editor and cannot be used to script standalone applications. In our work, we expose a similar breadth of functionality to Blueprints and the Unreal Editor, and our Python interface can be used to script standalone applications, live simulations running inside the editor, UE’s path tracer, and the editor itself, all through a unified and expressive high-level programming model.

3 Programming Model

In this section, we describe the design goals, fundamental concepts, and formal guarantees provided by our programming model (see Figures 3 and 4). In Section 4, we discuss our implementation of this model, including our low-level reflection interface and our method for passing NumPy arrays to and from UE. In our supplementary material, we include additional details about our model, including how we can use it to control the Unreal Editor and UE’s path tracer. **Assumptions** Throughout the following discussion, we assume the user’s Python code is being executed in a particular process on the *Python thread*. We further assume there is a UE application running in a different process that has been packaged with our SPEAR plugins, and that our plugins are executing the user’s UE work as eagerly as possible on the application’s *game thread*.




Fig. 5: Illustration of how the asynchronous operations in our programming model can be used to completely avoid blocking the UE game thread, thereby enabling a UE application to execute user Python code at native frame rates. Our Python client transforms a user’s Python code (top left) into a stream of commands (bottom left), and these commands are sent to our server running on its own thread within a UE application. We assume that the user’s Python code consists of only asynchronous operations, and therefore the server can avoid synchronizing with the game thread when queuing these operations (light grey boxes). When necessary, the server intentionally blocks (dotted colored boxes) to prevent the user’s Python code from getting too far ahead of the game thread. Once the game thread reaches the beginning of frame t , the system is in a steady state. In this steady state, as long as the user’s `begin_frame` and `end_frame` contexts execute faster than a UE frame, i.e., as long as the width of the sequence [`d`, `end`, `begin`, `e`, `execute`] on the server thread is less than the width of frame t on the game thread in the diagram above, then the UE application will execute at its native frame rate.

Design Goals When designing our programming model, we prioritize five main desiderata. *Expressive*. We want users to be able to express complex graphs of UE work, while also giving them precise control over how their graphs are executed within and across individual UE frames. *Programmable*. We want to expose as much UE functionality as possible, as directly as possible, without imposing opinionated domain concepts or abstractions on the user. *Ergonomic*. We want the syntax for calling UE functions and accessing UE variables to be as similar as possible to native Python. *Extensible*. We want it to be easy for users to expose new C++ functionality. *Efficient*. We want the speed of programs written in our model to be as close as possible to equivalent programs written directly in C++.

An important non-goal in our model is narrowing the expressive power of UE to make it easier to use for practitioners within a particular subdomain. There is no doubt that such interfaces are useful, but we believe they are best implemented as optional abstractions on top of our programming model, rather than as required abstractions within our model.

Specifying Graphs of Unreal Engine Work In our programming model, graphs of UE work are specified as transactions (see Figure 3). In particular, the user specifies a transaction by defining a `begin_frame` context followed by an `end_frame` context. Within each context, the user specifies a graph of UE work simply by implementing the graph as Python code. The UE work specified in each `begin_frame` context is guaranteed to execute at the beginning of a UE frame, the UE work specified in its corresponding `end_frame` context is guaranteed to execute at the end of the *same frame*, and all UE work is guaranteed to execute sequentially on the game thread in the order it is specified in Python.

By default, our programming model guarantees that each transaction executes within a single frame, but we intentionally do not guarantee that consec-



Fig. 6: We demonstrate the flexibility of SPEAR by using it to programmatically manipulate the ELECTRICDREAMS sample project from Epic Games. **(a):** We control UE’s procedural content generation (PCG) system by translating the main PCG entity in this scene (the rock structure in the center of each image) from left to right. Notice how the rock structure automatically harmonizes with the rest of the scene in a convincing way (e.g., the water adjusts around the rock, logs appear and connect with nearby structures), even when it is being driven by our simple programmatic control. **(b):** We simulate time-of-day changes by controlling the orientation of the scene’s sky light.

utive transactions execute on consecutive frames. In order to implement deterministic stepping in the absence of such a guarantee, the user can begin their UE simulation in a paused state, and in each subsequent transaction unpause the simulation, mutate the game state, and pause the simulation again. A practical advantage of this approach, as opposed to blocking the game thread between transactions (e.g., CARLA’s synchronous mode [18]), is that the UE application remains responsive to user input at all times. This interactivity is especially useful when controlling simulations running inside the Unreal Editor, because the editor can be used as a powerful visual debugging tool to navigate around the simulation environment and inspect the game state, even in-between transactions when the simulation is paused. Alternatively, if desired, the user can explicitly prevent the game thread from advancing in-between transactions by calling `end_frame(single_step=True)`, which instructs the game thread to advance to the beginning of the next frame and wait for the next transaction.

Calling Functions and Accessing Variables Given a Python object that represents an underlying UE object (e.g., `bp_axes` in Figure 3), any function or variable on the underlying UE object that is visible to the reflection system can be accessed as though it was a native Python function or attribute on the Python object (e.g., calling the `SetActorScale3D` function and accessing the `RootComponent` variable in Figure 3). Our programming model also supports accessing variables that are nested arbitrarily deeply within structs, e.g., `bp_axes.RootComponent.RelativeLocation.X.get()` would be a valid expression in Figure 3.



Fig. 7: SPEAR can be used in co-simulation applications with external physics simulators. In this application, we interactively control the MuJoCo physics simulator [72] using the default MuJoCo viewer, e.g., by applying a force to the leftmost chair (red arrow). In real-time as the MuJoCo simulation is running, we query the state of the MuJoCo scene (inset images), and we use SPEAR to update the state of a corresponding UE scene (large images).

Exposing new Functions and Variables It is trivial to expose new C++ functions and variables to UE’s reflection system, and therefore to SPEAR, simply by adding a `UFUNCTION` or `UPROPERTY` annotation next to the function or variable (e.g., `MyFunction` or `MyProperty` in Figure 3) in any C++ header file, including in header files that are outside of the SPEAR codebase. An important consequence of this design is that new C++ functionality can be exposed to Python without modifying SPEAR code, which is not possible in any existing UE-based simulator. No additional registration steps or code wrappers are required to expose new C++ functionality to Python.

Synchronizing with the Game Thread By default, every UE operation in our programming model (e.g., calling a UE function, accessing a UE property) is fully synchronous. In other words, each operation is guaranteed to finish executing on the game thread before control is returned to the user’s Python code. However, for improved efficiency, we also provide an asynchronous variant for each UE operation in our model (e.g., `async.GetComponentLocation` in Figure 3), which avoids synchronizing with UE and immediately returns a future object. This future object can be used later in the user’s Python program to obtain the return value of the underlying operation (e.g., `location` in Figure 3). Each asynchronous operation is guaranteed to execute on the game thread in the same order it would have if the user invoked its synchronous counterpart, and getting the return value from an asynchronous operation will not synchronize with the game thread unless the operation is still pending, in which case getting its return value will block until the operation is complete.

When specifying transactions using asynchronous operations, it is possible for the Python thread to execute significantly faster than the game thread. To prevent these two threads from diverging, we intentionally gate the progress of the Python thread by allowing at most one pending transaction at a time. Specifically, if work from the user’s previous `begin_frame` context is still pending, then the next attempt to enter a `begin_frame` context will block until the previous work has finished, and likewise for `end_frame`.

Executing User Python Code at Native Frame Rates An important feature of our programming model is that asynchronous operations can be used to completely avoid blocking the game thread, thereby enabling UE to execute user Python code at native frame rates (see Figure 5).



Fig. 8: We demonstrate the flexibility of our camera sensor by using it to render synchronized multi-view images of a detailed human character in the META HUMANS sample project from Epic Games.

Summary Our programming model is *expressive* because it enables users to specify complex graphs of UE work simply by writing Python code, and because our approach for specifying graphs as transactions gives users precise control over how their graphs are executed. Our model is *programmable* and *ergonomic* because we expose all reflection-visible UE functionality as native Python functions and attributes. Our model is *extensible* because it is trivial to expose new C++ functions and variables to the UE reflection system, and therefore to SPEAR. Finally, our model is *efficient* because it can be used to completely avoid blocking the game thread, thereby enabling UE to execute user Python code at native frame rates.

4 System Architecture

We implement our programming model using a client-server architecture that is conceptually similar to those found in existing simulators [16, 18, 21, 28, 36, 47, 48, 58, 67, 87], but with some unique characteristics that are needed for performance and to deliver on the guarantees provided by our programming model. In our supplementary material, we include additional details about our system architecture, including our high-level Python interface and the implementation of our camera sensor.

Client-Server Interface We implement a Python client that is responsible for transforming user-defined transactions of UE work into streams of commands that are sent to the server. Additionally, we implement a C++ server that runs as a plugin within the UE application and responds to commands from the client. The client communicates with the server over a TCP/IP connection, and it is therefore possible for the client and server to run on different machines.

We implement both sides of our client-server interface in C++ using `rpclib` [70], and we implement a Python wrapper for the client side of our interface using `nanobind` [26]. This design enables us to implement the server side of our interface as a collection of strongly typed C++ entry points that permit standard containers and our own custom types as arguments and return values. Additionally, this design enables our client to call server entry points as though they were native C++ functions. Throughout this discussion, when we say informally that the client sends a command to the server, we mean precisely that the client synchronously calls one of our strongly typed entry points and obtains a strongly typed return value.



Fig. 9: We demonstrate the flexibility of SPEAR by using it to implement an agentic natural-language scene editing application, where a vision-and-language coding assistant manipulates a UE scene by iteratively writing SPEAR programs in response to user text prompts.

The Server Thread In order to support the asynchronous operations in our programming model, we run our server on a separate thread within the UE application, which we refer to as the *server thread*. We emphasize that the asynchronous operations in our model are asynchronous with respect to the server thread and the game thread, but they are synchronous with respect to the Python thread and the server thread. This setup simplifies our system architecture, while also ensuring that our server can respond to commands without needing to synchronize with the game thread, in contrast to existing UE plugins (e.g., [58, 87]).

Because our server runs on its own thread, it needs a well-defined way of interacting with the game thread, since nearly all UE work specified by the client must access the game thread at some point. With this goal in mind, we implement a thread-safe task queuing system. In particular, we maintain two separate thread-safe queues for `begin_frame` and `end_frame` work respectively. Whenever an entry point is called that needs to access the game thread, the server places a task in one of these queues, and the queues are drained by the game thread at the beginning and end of each UE frame in a way that respects the execution ordering guarantees in our programming model.

Synchronizing Progress Across Threads In our system architecture, we rely on the server to synchronize the progress of both the Python thread and the game thread according to the rules of our programming model. To perform this synchronization, the server must track the state of the user’s Python code as it enters and exits `begin_frame` and `end_frame` contexts.

We keep the server up-to-date by implementing the following logic in our `begin_frame` and `end_frame` context managers [63] on the client. When the user’s Python code enters a `begin_frame` context, the client sends a `begin` command to the server; when the user’s Python code leaves a `begin_frame` context, the client sends an `execute` command to the server; and when the user’s Python code leaves an `end_frame` context, the client sends an `end` command to the server. The `begin` and `execute` commands will block if the server determines that the user’s Python code is too far ahead of the game thread (see Figure 5).

Hand-Crafted Server Entry Points We implement 193 hand-crafted server entry points to expose various UE functions that are not visible to the reflection system, as well as the reflection system itself. We implement each synchronous variant manually, and we generate its asynchronous variant automatically using template metaprogramming [74]. This technique leverages the fact that each syn-

chronous variant is implemented as a C++ function, rather than as a text-based command that must be parsed and dispatched imperatively (e.g., as in [58, 87]). When the user’s Python code is inside a `begin_frame` or `end_frame` context, all of our server entry points can be called as Python functions (e.g., `load_class` and `spawn_actor` in Figure 3). Internally, our Python client routes each such function call through our nanobind wrapper and ultimately to the appropriate entry point on the server.

Exposing the Reflection System Roughly 75% of our hand-crafted server entry points are dedicated to exposing some aspect of the UE reflection system. Together, these entry points enable our Python client to find classes, call functions, and manipulate variables on objects dynamically at runtime using strings as keys, i.e., without requiring hand-crafted wrapper code for each class, function, and variable. Through this approach, we ultimately expose over 14K UE functions and 53K UE properties to Python.

When exposing UE’s reflection system to Python, we leverage the fact that UE imposes some restrictions on the types that it considers to be *reflectable* [9]. In other words, only certain types are allowed in reflected function signatures and member variables (e.g., primitive types, strings, pointers to UE objects, some containers, `enums`, and `structs` that are recursively composed of types from this set). These restrictions enable UE to automatically serialize and deserialize to and from JSON any variable whose type is reflectable, which is especially useful for our purposes, because Python dictionaries can also be automatically serialized and deserialized to and from JSON. The structural alignment between reflectable types in UE and dictionaries in Python enables us to use dictionaries as a universal representation for getting and setting variables, as well as passing arguments and return values to and from functions in user Python code (e.g., `NewScale3D` in Figure 3).

NumPy Interoperability When passing large blocks of data (e.g., high-resolution images) to and from UE functions, our approach for representing arguments and return values as JSON strings can be inefficient. We address this limitation by implementing a dispatching system for invoking custom functions on UE objects, which supports NumPy arrays as arguments and return values and marshals them efficiently. In our system, any UE object that has a hierarchy of child components (see [4] for a more detailed discussion of UE’s component system) can define custom functions by inserting a special child component into its hierarchy and binding named functions to this component at runtime. We refer to these named functions as `SPFUNCTIONS`, and they can be called from user Python code as though they were reflection-visible UE functions, i.e., they appear to user Python code as native Python functions in our programming model. Defining an `SPFUNCTION` in C++ code is trivial, requiring a similar amount of effort as exposing a function to UE’s reflection system (see our supplementary material for a code example).

For simplicity, we restrict the signature of `SPFUNCTIONS` in the following way. All `SPFUNCTIONS` must take as input, and return as output, three distinct objects: (1) a collection of named data arrays with shape and type metadata;

Simulator	Functions	Variables	UE functions	UE variables	Lines of code
AirSim	92	189	0	0	144,536
CARLA	465	508	0	0	150,502
UnrealCV+	56	0	747	8,721	11,301
SPEAR (ours)	193	67	14,485	53,537	27,193

Table 1: Comparison of programmable functionality across UE simulators. We report the number of hand-crafted functions and variables each simulator [18, 58, 67, 87] provides, as well as the number of underlying UE functions and variables it exposes (higher is better). For reference, we also report the total lines of Python and C++ code in each simulator codebase (lower is better). SPEAR provides an order of magnitude more programmable functionality than existing UE-based simulators while maintaining a modest code footprint.

(2) a collection of named UE objects represented as JSON strings; and (3) a string that can be used to encode miscellaneous user data. When a call to an SPFUNCTION reaches our client-server boundary, any NumPy arrays that are given as arguments are mapped to our named data array representation, and any data arrays that are returned as output are mapped back to NumPy arrays.

Interprocess Shared Memory We further optimize the efficiency of our system using interprocess shared memory [66], or simply *shared memory*. Using shared memory, we can, e.g., transfer rendered images from the GPU directly into a user’s NumPy array with no additional data copying.

We integrate shared memory into our system by implementing a C++ interface for allocating and deallocating shared memory regions, and we expose this interface to Python through a set of hand-crafted server entry points. Owing to this interface, user Python code can optionally allocate a shared memory region, create a NumPy array that is backed by that region, and pass the NumPy array to an SPFUNCTION. Importantly, when our server calls an SPFUNCTION implementation, we internally resolve references to shared memory for all arguments prior to calling the implementation. As a result, the implementation remains completely decoupled from the caller’s decision to use shared memory for arguments. We implement a similar strategy when returning data from an SPFUNCTION, so the user’s Python code remains decoupled from the implementation’s decision to use shared memory for return values.

5 Results

Example Applications We demonstrate the utility of SPEAR through a diverse collection of example applications (see Figures 1, 6, 7, 8, and 9). We include several additional applications in our supplementary material, and all of our applications are included in our public code repository.

Every example application we show in our figures and videos would not be possible to implement in any existing UE-based simulator because existing simulators do not expose the necessary functionality (e.g., to interact with UE’s procedural content system, or control UE’s path tracer, etc). Additionally, nearly

Configuration	Time (ms)	FPS	
Standalone	7.7	129.9	
Standalone + extra work	17.7	56.5	
UnrealCV+	286.9	3.5	
SPEAR (ours)			
× async × shared mem	40.5	24.7	
× async ✓ shared mem	31.6	31.7	
✓ async × shared mem	37.3	26.8	
✓ async ✓ shared mem	17.8	56.2	
SPEAR (ours)			
1 frame rendering latency	15.4	64.8	
2 frame rendering latency	13.6	73.4	

Table 2: Comparison of rendering performance under various configurations when rendering the image on the right at 1920×1080 resolution. For each configuration, we measure the total end-to-end time required to deliver a rendered image to a user’s Python program, and we report frame time (ms) and frames per second (FPS). As baselines, we include a standalone UE executable that does not communicate with Python, a standalone executable that does not communicate with Python but does the extra UE work necessary to do so (e.g., rendering an extra view of the scene to an off-screen buffer), and UnrealCV+ [58,87]. We report the performance of SPEAR with and without asynchronous communication and shared memory, and with different amounts of rendering latency. We find that asynchronous communication and shared memory both improve performance, as does increasing rendering latency, and that SPEAR is 9–21 $\times$ faster than UnrealCV+ while also supporting more photorealistic rendering.

all of the functionality required in these applications is already exposed through UE’s reflection system, and is therefore exposed automatically in SPEAR. In summary, SPEAR exposes more functions than existing simulators (i.e., SPEAR is more programmable), and this makes it possible to implement a wider range of programs than would be possible otherwise. See our supplementary material for a more detailed discussion.

Programmable functionality In Table 1, we compare programmable functionality and code size across UE-based simulators. For each baseline simulator, we include all publicly documented Python functions in our count of its hand-crafted functions. If a simulator can be configured via some kind of global settings file, we include all documented settings in our count of its hand-crafted variables. UnrealCV+ [58,87] uses a text-based command interface, so we include all publicly documented text commands in our count of its hand-crafted functions. Both UnrealCV+ and SPEAR enable users to execute UE console commands and access UE console variables [2], so we include these functions and variables in our counts for both simulators. For SPEAR, we report the number of unique server entry points and YAML configuration options as our count of hand-crafted functions and variables. To count the number of UE functions and variables we expose in SPEAR, we use our reflection interface to iterate over all reflection-visible UE functions and variables (i.e., all `UFUNCTIONS` and `UPROPERTIES`).

Simulator	Time (ms)	FPS
0 frames rendering latency		
AirSim	379.4	2.6
SPEAR (ours)	31.0	32.3
2 frames rendering latency		
CARLA	30.6	32.7
SPEAR (ours)	27.0	37.1



Table 3: Comparison of rendering performance across simulators when rendering the images on the right at 1920×1080 resolution. For each simulator [18, 67], we measure the total end-to-end time required to deliver a rendered image to a user’s Python program, and we report frame time (ms) and frames per second (FPS). We ensure the rendering speed of each simulator is as similar as possible when running as a standalone executable that does not communicate with Python (SPEAR: 89 FPS; CARLA: 90 FPS; AirSim: 93 FPS). Under these conditions, any differences we observe when measuring total end-to-end rendering time are attributable to differences in communication overhead. We find that SPEAR is $12\times$ faster than AirSim, and is 10% faster than CARLA under matched amounts of rendering latency.

Rendering Performance We perform all performance comparisons on a Windows 11 desktop workstation with an NVIDIA GeForce RTX 4090 GPU, a 4.5 GHz AMD Ryzen 9 processor with 16 physical cores, and 192 GB of memory. We include additional performance comparisons in our supplementary material.

In Table 2, we compare rendering performance relative to UnrealCV+ [58, 87] and a standalone executable that does not communicate with Python. For this experiment, we build one executable that includes our SPEAR plugins, and another that includes the UnrealCV+ plugin, in an otherwise identical UE project. This approach enables us to evaluate the rendering performance of both plugins under precisely controlled conditions (e.g., same view of the same scene with all of the same project settings). The SPEAR camera sensor is designed to incur a user-configurable amount of rendering latency in exchange for increased throughput (see our supplementary material for details). All rows in Table 2 are configured to incur 0 frames of latency unless otherwise noted, and all rows that incur 1 or 2 frames of rendering latency use shared memory.

In Table 3, we compare rendering performance across UE-based simulators. For this experiment, it is not straightforward to render the same scene in all simulators, because AirSim does not distribute any non-trivial environments in source form, and CARLA environments are tightly coupled to the CARLA source code. So we render a different scene in each simulator, but we ensure the rendering speed of each simulator is as similar as possible when running as a standalone executable that does not communicate with Python (SPEAR: 89 FPS; CARLA: 90 FPS; AirSim: 93 FPS). Under these conditions, any differences we observe when measuring total end-to-end rendering time with Python are attributable to differences in communication overhead.

For AirSim and CARLA, we use their default environments, and we select views of each environment that render at roughly the same speed. For SPEAR, we use the HILLSIDESAMPLE project from Epic Games, because it depicts a

city-scale urban outdoor environment that is qualitatively similar to the default environments in AirSim and CARLA. However, HILLSIDESAMPLE is designed to be a challenging next-generation showcase for UE’s latest rendering features, and therefore it renders noticeably slower than the default AirSim and CARLA environments. To account for this discrepancy, we remove distant background objects and adjust rendering quality settings until we achieve roughly the same standalone rendering speed across all simulators (see above). AirSim and CARLA incur 0 and 2 frames of rendering latency respectively, so we report the performance of SPEAR with each of these latency settings to ensure fair comparisons.

6 Conclusions

In this work, we introduced SPEAR: A Simulator for Photorealistic Embodied AI Research. At its core, SPEAR is a Python library that can programmatically control any Unreal Engine (UE) application using an expressive high-level programming model. The SPEAR programming model exposes over 14K unique UE functions to Python (an order of magnitude more programmable functionality than existing UE-based simulators), and enables photorealistic rendering at over 150 megapixels per second (an order of magnitude faster than existing UE plugins). Additionally, the default SPEAR camera sensor is highly configurable, and provides ground truth image modalities that are not available in any existing UE-based simulator. We demonstrated the utility of SPEAR through a diverse collection of example applications, including controlling multiple embodied agents in a photorealistic city-scale environment, controlling UE’s procedural content generation systems, rendering multi-view images of detailed human faces, coordinating an interactive co-simulation with MuJoCo, and editing scenes with natural language via an AI coding assistant.

We believe that SPEAR could become a foundational data engine in computer vision, robotics, and embodied AI. In the near future, SPEAR could be used to train agile robots in city-scale photorealistic environments, as well as interactive world models that understand the dynamics and spatial structure of the physical world. By providing programmability via small Python programs, and observability via fast photorealistic rendering, SPEAR is ideally positioned as a bridge to connect internet-scale vision-and-language models to state-of-the-art virtual worlds in the Unreal Engine. In turn, this could lead to new forms of AI-assisted content creation and personalized entertainment, as well as new virtual laboratories for studying the foundations of spatial intelligence.

Acknowledgments

We thank Ahlam Laatiki for creating the default SPEAR apartment scene; Epic Games for making the sample projects that we use throughout this paper available for download; Songyou Peng for his excellent guidance on crafting our rebuttal; and Kevin Blackburn-Matzen, Samarth Brahmhatt, Martí Ferragut Galtes, and David Hafner for their assistance with early prototyping.

References

1. Blueprints visual scripting. Epic Games Developer Documentation 5
2. Console variables and commands. Epic Games Developer Documentation 13
3. NVIDIA Isaac Sim. <https://developer.nvidia.com/isaac/sim> 2
4. An overview of components. Epic Games Developer Documentation 11
5. Panda3D. <https://www.panda3d.org> 2
6. Scripting the Unreal Editor using Python. Epic Games Developer Documentation 5
7. Unity. <https://unity.com> 2
8. Unreal Engine. <https://www.unrealengine.com> 2
9. Unreal Engine reflection system. Epic Games Developer Documentation 3, 11
10. Black Myth: Wukong tops 25 million copies sold as its merchandise sales skyrocket in China. Game World Observer (2025) 2
11. Clair Obscur takes home 9 Game Awards including Game of the Year. Epic Games Store News (2025) 2
12. Bachmann, R., Kar, O.F., Mizrahi, D., Garjani, A., Gao, M., Griffiths, D., Hu, J., Dehghan, A., Zamir, A.: 4M-21: An any-to-any vision model for tens of tasks and modalities. In: NeurIPS 2024 2
13. Bai, J., Xia, M., Fu, X., Wang, X., Mu, L., Cao, J., Liu, Z., Hu, H., Bai, X., Wan, P., Zhang, D.: ReCamMaster: Camera-controlled generative rendering from a single video. In: ICCV 2025 2
14. Beattie, C., Leibo, J.Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., Petersen, S.: DeepMind Lab. arXiv 2016 4
15. Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: The arcade learning environment: An evaluation platform for general agents. Journal of Artificial Intelligence Research 47(1) (2013) 4
16. Bordes, F., Shekhar, S., Ibrahim, M., Bouchacourt, D., Vincent, P., Morcos, A.S.: PUG: Photorealistic and semantically controllable synthetic data for representation learning. In: NeurIPS 2023 2, 9
17. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: OpenAI Gym. arXiv 2016 5
18. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: CoRL 2017 1, 2, 5, 7, 9, 12, 14
19. Fan, L., Wang, G., Jiang, Y., Mandlekar, A., Yang, Y., Zhu, H., Tang, A., Huang, D.A., Zhu, Y., Anandkumar, A.: MineDojo: Building open-ended embodied agents with internet-scale knowledge. In: NeurIPS 2022 Datasets and Benchmarks Track 4
20. Gabeur, V., Long, S., Peng, S., Voigtlaender, P., Sun, S., Bao, Y., Truong, K., Wang, Z., Zhou, W., Barron, J.T., Genova, K., Kannen, N., Ben, S., Li, Y., Guo, M., Yogin, S., Gu, Y., Chen, H., Wang, O., Xie, S., Zhou, H., He, K., Funkhouser, T., Alayrac, J.B., Soricut, R.: Image generators are generalist vision learners. arXiv 2026 2
21. Gan, C., Schwartz, J., Alter, S., Mrowca, D., Schrimpf, M., Traer, J., Freitas, J.D., Kubilius, J., Bhandwaldar, A., Haber, N., Sano, M., Kim, K., Wang, E., Lingelbach, M., Curtis, A., Feiglis, K., Bear, D.M., Gutfreund, D., Cox, D., Torralba, A., DiCarlo, J.J., Tenenbaum, J.B., McDermott, J.H., Yamins, D.L.K.: ThreeD-World: A platform for interactive multi-modal physical simulation. In: NeurIPS 2021 Datasets and Benchmarks Track 2, 9

22. Ge, Y., Tang, Y., Xu, J., Gokmen, C., Li, C., Ai, W., Martinez, B.J., Aydin, A., Anvari, M., Chakravarthy, A.K., Yu, H.X., Wong, J., Srivastava, S., Lee, S., Zha, S., Itti, L., Li, Y., Martin-Martin, R., Liu, M., Zhang, P., Zhang, R., Fei-Fei, L., Wu, J.: BEHAVIOR vision suite: Customizable dataset generation via simulation. In: CVPR 2024 [2](#)
23. Gong, R., Huang, J., Zhao, Y., Geng, H., Gao, X., Wu, Q., Ai, W., Zhou, Z., Terzopoulos, D., Zhu, S.C., Jia, B., Huang, S.: Arnold: A benchmark for language-grounded task learning with continuous states in realistic 3D scenes. In: ICCV 2023 [2](#)
24. Hwangbo, J., Lee, J., Dosovitskiy, A., Bellicoso, D., Tsounis, V., Koltun, V., Hutter, M.: Learning agile and dynamic motor skills for legged robots. *Science Robotics* **4**(26) (2019) [2](#)
25. Jaderberg, M., Czarnecki, W.M., Dunning, I., Marris, L., Lever, G., Castañeda, A.G., Beattie, C., Rabinowitz, N.C., Morcos, A.S., Ruderman, A., Sonnerat, N., Green, T., Deason, L., Leibo, J.Z., Silver, D., Hassabis, D., Kavukcuoglu, K., Graepel, T.: Human-level performance in 3D multiplayer games with population-based deep reinforcement learning. *Science* **364**(6443) (2019) [4](#)
26. Jakob, W.: nanobind: tiny and efficient C++/Python bindings. <https://github.com/wjakob/nanobind> [9](#)
27. Johnson, M., Hofmann, K., Hutton, T., Bignell, D.: The Malmo platform for artificial intelligence experimentation. In: IJCAI 2016 [4](#)
28. Juliani, A., Berges, V.P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., Lange, D.: Unity: A general platform for intelligent agents. *arXiv* 2018 [2](#), [9](#)
29. Karis, B., Stubbe, R., Wihlidal, G.: A deep dive into Nanite virtualized geometry. In: SIGGRAPH 2021 Course on Advances in Real-Time Rendering and Games [2](#)
30. Kaufmann, E., Bauersfeld, L., Loquercio, A., Müller, M., Koltun, V., Scaramuzza, D.: Champion-level drone racing using deep reinforcement learning. *Nature* **620**(7976) (2023) [1](#)
31. Kaufmann, E., Loquercio, A., Ranftl, R., Dosovitskiy, A., Koltun, V., Scaramuzza, D.: Deep drone racing: Learning agile flight in dynamic environments. In: CoRL 2018 [1](#)
32. Kaufmann, E., Loquercio, A., Ranftl, R., Müller, M., Koltun, V., Scaramuzza, D.: Deep drone acrobatics. In: RSS 2020 [1](#)
33. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daut, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: CVPR 2024 [2](#)
34. Kempka, M., Wydmuch, M., Runc, G., Toczek, J., Jaśkowski, W.: ViZDoom: A Doom-based AI research platform for visual reinforcement learning. In: Computational Intelligence and Games 2016 [4](#)
35. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: ICCV 2023 [2](#)
36. Kolve, E., Mottaghi, R., Han, W., VanderBilt, E., Weihs, L., Herrasti, A., Gordon, D., Zhu, Y., Gupta, A., Farhadi, A.: AI2-THOR: An interactive 3D environment for visual AI. *arXiv* 2017 [2](#), [9](#)
37. Krähenbühl, P.: Free supervision from video games. In: CVPR 2018 [4](#)
38. Langmead, A.: Procedural content generation in UE5. In: Game Developers Conference 2023 [2](#)

39. Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., Hutter, M.: Learning quadrupedal locomotion over challenging terrain. *Science Robotics* **5**(47) (2020) [2](#)
40. Lerer, A., Gross, S., Fergus, R.: Learning physical intuition of block towers by example. In: *ICML 2016* [2](#)
41. Li, C., Xia, F., Martín-Martín, R., Lingelbach, M., Srivastava, S., Shen, B., Vainio, K.E., Gokmen, C., Dharan, G., Jain, T., Kurenkov, A., Liu, C.K., Gweon, H., Wu, J., Fei-Fei, L., Savarese, S.: iGibson 2.0: Object-centric simulation for robot learning of everyday household tasks. In: *CoRL 2021* [4](#)
42. Li, C., Zhang, R., Wong, J., Gokmen, C., Srivastava, S., Martín-Martín, R., Wang, C., Levine, G., Ai, W., Martinez, B., Yin, H., Lingelbach, M., Hwang, M., Hiranaka, A., Garlanka, S., Aydin, A., Lee, S., Sun, J., Anvari, M., Sharma, M., Bansal, D., Hunter, S., Kim, K.Y., Lou, A., Matthews, C.R., Villa-Renteria, I., Tang, J.H., Tang, C., Xia, F., Li, Y., Savarese, S., Gweon, H., Liu, C.K., Wu, J., Fei-Fei, L.: BEHAVIOR-1K: A human-centered, embodied AI benchmark with 1,000 everyday activities and realistic simulation. In: *CoRL 2022* [2](#)
43. Loquercio, A., Kaufmann, E., Ranftl, R., Dosovitskiy, A., Koltun, V., Scaramuzza, D.: Deep drone racing: From simulation to reality with domain randomization. *Transactions on Robotics* **36**(1) (2020) [1](#)
44. Loquercio, A., Kaufmann, E., Ranftl, R., Müller, M., Koltun, V., Scaramuzza, D.: Learning high-speed flight in the wild. *Science Robotics* **6**(59) (2021) [1](#)
45. Maes, P.: Concepts and experiments in computational reflection. In: *OOPSLA 1987* [3](#)
46. Magne, L., Awadalla, A., Wang, G., Xu, Y., Belofsky, J., Hu, F., Kim, J., Schmidt, L., Gkioxari, G., Kautz, J., Yue, Y., Choi, Y., Zhu, Y., Fan, L.J.: NitroGen: An open foundation model for generalist gaming agents. *arXiv* 2026 [4](#)
47. Martinez-Gonzalez, P., Oprea, S., Castro-Vargas, J.A., Garcia-Garcia, A., Orts-Escolano, S., Garcia-Rodriguez, J., Vincze, M.: UnrealROX+: An improved tool for acquiring synthetic data from virtual 3D environments. In: *IJCNN 2021* [2](#), [9](#)
48. Martinez-Gonzalez, P., Oprea, S., Garcia-Garcia, A., Jover-Alvarez, A., Orts-Escolano, S., Garcia-Rodriguez, J.: UnrealROX: An extremely photorealistic virtual reality environment for robotics simulations and synthetic data generation. *Virtual Reality* **24** (2019) [2](#), [9](#)
49. Mazur, K., Bae, G., Davison, A.J.: SuperPrimitive: Scene reconstruction at a primitive level. In: *CVPR 2024* [2](#)
50. Miki, T., Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., Hutter, M.: Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics* **7**(62) (2022) [2](#)
51. Mizrahi, D., Bachmann, R., Kar, O.F., Yeo, T., Gao, M., Dehghan, A., Zamir, A.: 4M: Massively multimodal masked modeling. In: *NeurIPS 2023* [2](#)
52. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. *Nature* **518**(7540) (2015) [4](#)
53. Müller, M., Dosovitskiy, A., Ghanem, B., Koltun, V.: Driving policy transfer via modularity and abstraction. In: *CoRL 2018* [1](#)
54. Narkowicz, K., Costa, T.: MegaLights: Stochastic direct lighting in Unreal Engine 5. In: *SIGGRAPH 2025 Advances in Real-Time Rendering in Games* [2](#)

55. OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., Zhang, L.: Solving Rubik’s Cube with a robot hand. arXiv 2019 [2](#)
56. OpenAI, Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., d. O. Pinto, H.P., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., Zhang, S.: Dota 2 with large scale deep reinforcement learning. arXiv 2019 [1](#), [4](#)
57. Puig, X., Undersander, E., Szot, A., Cote, M.D., Yang, T.Y., Partsey, R., Desai, R., Clegg, A.W., Hlavac, M., Min, S.Y., Vondruš, V., Gervet, T., Berges, V.P., Turner, J.M., Maksymets, O., Kira, Z., Kalakrishnan, M., Malik, J., Chaplot, D.S., Jain, U., Batra, D., Rai, A., Mottaghi, R.: Habitat 3.0: A co-habitat for humans, avatars and robots. In: ICLR 2024 [4](#)
58. Qiu, W., Zhong, F., Zhang, Y., Qiao, S., Xiao, Z., Kim, T.S., Wang, Y.: UnrealCV: Virtual worlds for computer vision. In: Multimedia 2017 [2](#), [5](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
59. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Alwala, K.V., Carion, N., Wu, C.Y., Girshick, R., Dollár, P., Feichtenhofer, C.: SAM 2: Segment anything in images and videos. In: ICLR 2025 [2](#)
60. Richter, S., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV 2016 [1](#), [4](#)
61. Richter, S.R., Hayder, Z., Koltun, V.: Playing for benchmarks. In: ICCV 2017 [1](#), [4](#)
62. Roberts, M., Ramapuram, J., Ranjan, A., Kumar, A., Bautista, M.A., Paczan, N., Webb, R., Susskind, J.M.: Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In: ICCV 2021 [2](#), [3](#)
63. van Rossum, G., Coghlan, A.: PEP 343: The “with” statement. Python Enhancement Proposals [10](#)
64. Savva, M., Chang, A.X., Dosovitskiy, A., Funkhouser, T., Koltun, V.: MINOS: Multimodal indoor simulator for navigation in complex environments. arXiv 2017 [4](#)
65. Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., Parikh, D., Batra, D.: Habitat: A platform for embodied AI research. In: ICCV 2019 [4](#)
66. Schäling, B.: The Boost C++ Libraries. XML Press (2011) [12](#)
67. Shah, S., Dey, D., Lovett, C., Kapoor, A.: AirSim: High-fidelity visual and physical simulation for autonomous vehicles. In: Field and Service Robotics 2017 [2](#), [5](#), [9](#), [12](#), [14](#)
68. Shen, B., Xia, F., Li, C., Martín-Martín, R., Fan, L., Wang, G., Pérez-D’Arpino, C., Buch, S., Srivastava, S., Tchapmi, L.P., Tchapmi, M.E., Vainio, K., Wong, J., Fei-Fei, L., Savarese, S.: iGibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In: IROS 2021 [4](#)
69. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3. arXiv 2025 [2](#)
70. Szelei, T.: rpclib: Modern msgpack-rpc for C++. <http://rpclib.net> [9](#)
71. Szot, A., Clegg, A., Undersander, E., Wijmans, E., Zhao, Y., Turner, J., Maestre, N., Mukadam, M., Chaplot, D., Maksymets, O., Gokaslan, A., Vondrus, V., Dharur,

- S., Meier, F., Galuba, W., Chang, A., Kira, Z., Koltun, V., Malik, J., Savva, M., Batra, D.: Habitat 2.0: Training home assistants to rearrange their habitat. In: NeurIPS 2021 [4](#), [5](#)
72. Todorov, E., Erez, T., Tassa, Y.: MuJoCo: A physics engine for model-based control. In: IROS 2012 [3](#), [5](#), [8](#)
73. Tong, S., Brown, E., Wu, P., Woo, S., Middepogu, M., Akula, S.C., Yang, J., Yang, S., Iyer, A., Pan, X., Wang, Z., Fergus, R., LeCun, Y., Xie, S.: Cambrian-1: A fully open, vision-centric exploration of multimodal LLMs. In: NeurIPS 2024 [2](#)
74. Vandevoorde, D., Josuttis, N.M., Gregor, D.: C++ Templates: The Complete Guide. Addison-Wesley Professional, 2 edn. (2017) [10](#)
75. Vinyals, O., Babuschkin, I., Czarnecki, W.M., Mathieu, M., Dudzik, A., Chung, J., Choi, D.H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J.P., Jaderberg, M., Vezhnevets, A.S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T.L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., Silver, D.: Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575** (2019) [1](#), [4](#)
76. Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhnevets, A.S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., Quan, J., Gaffney, S., Petersen, S., Simonyan, K., Schaul, T., van Hasselt, H., Silver, D., Lillicrap, T., Calderone, K., Keet, P., Brunasso, A., Lawrence, D., Ekermo, A., Repp, J., Tsing, R.: StarCraft II: A new challenge for reinforcement learning. *arXiv* 2017 [1](#), [4](#)
77. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: VGGT: Visual geometry grounded transformer. In: CVPR 2025 [2](#)
78. Wang, J., Chen, M., Zhang, S., Karaev, N., Schönberger, J., Labatut, P., Bojanowski, P., Novotny, D., Vedaldi, A., Rupprecht, C.: VGGT-2. In: CVPR 2026 [2](#)
79. Wright, D., Narkowicz, K., Kelly, P.: Lumen: Real-time global illumination in Unreal Engine 5. In: SIGGRAPH 2022 Advances in Real-Time Rendering in Games [2](#)
80. Wu, W., He, H., He, J., Wang, Y., Duan, C., Liu, Z., Li, Q., Zhou, B.: MetaUrban: An embodied AI simulation platform for urban micromobility. In: ICLR 2025 [2](#)
81. Wu, W., He, H., Zhang, C., He, J., Zhao, S.Z., Gong, R., Li, Q., Zhou, B.: Towards autonomous micromobility through scalable urban simulation. In: CVPR 2025 [2](#)
82. Wurman, P.R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T.J., Capobianco, R., Devlic, A., Eckert, F., Fuchs, F., Gilpin, L., Khandelwal, P., Kompella, V., Lin, H., MacAlpine, P., Oller, D., Seno, T., Sherstan, C., Thomure, M.D., Aghabozorgi, H., Barrett, L., Douglas, R., Whitehead, D., Dürr, P., Stone, P., Spranger, M., Kitano, H.: Outracing champion Gran Turismo drivers with deep reinforcement learning. *Nature* **602** (2022) [1](#), [4](#)
83. Xia, F., Zamir, A.R., He, Z.Y., Sax, A., Malik, J., Savarese, S.: Gibson Env: Real-world perception for embodied agents. In: CVPR 2018 [4](#)
84. Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A.X., Guibas, L.J., Su, H.: SAPIEN: A simulated part-based interactive environment. In: CVPR 2020 [4](#)
85. Yang, S., Yang, J., Huang, P., Brown, E., Yang, Z., Yu, Y., Tong, S., Zheng, Z., Xu, Y., Wang, M., Lu, D., Fergus, R., LeCun, Y., Fei-Fei, L., Xie, S.: Cambrian-S: Towards spatial supersensing in video. *arXiv* 2025 [2](#)

- 86. Ye, X., Ren, J., Zhuang, Y., He, X., Liang, Y., Yang, Y., Zhong, X., Dogra, M., Liu, E., Benavente, K., Nagaraju, R.M., Sharma, D.V., Ma, Z., Shu, T., Hu, Z., Qin, L.: SimWorld: An open-ended simulator for agents in physical and social worlds. In: NeurIPS 2025 [2](#)
- 87. Zhong, F., Wu, K., Wang, C., Chen, H., Ci, H., Li, Z., Wang, Y.: UnrealZoo: Enriching photo-realistic virtual worlds for embodied AI. In: ICCV 2025 [2](#), [5](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
- 88. Zhou, B., Krähenbühl, P., Koltun, V.: Does computer vision matter for action? Science Robotics **4**(30) (2018) [1](#), [4](#)