

# Rapid Development of Compositional AI

Lee Martie, Jessie Rosenberg, Veronique Demers, Gaoyuan Zhang, Onkar Bhardwaj, John Henning, Aditya Prasad, Matt Stallone, Ja Young Lee, Lucy Yip, Damilola Adesina, Elahe Paikari, Oscar Resendiz, Sarah Shaw, David Cox  
MIT-IBM Watson AI Lab, IBM Research  
Cambridge, Massachusetts 02142 USA

{lee.martie, jrosenb, vdemers, gaoyuan.zhang, onkarbhardwaj, john.l.henning, aditya.prasad, mstallone, ja.young.lee, lucy.yip, dami.adesina, epaikari, oscar.resendiz, sarah.y.shaw, david.d.cox}@ibm.com

**Abstract**—Compositional AI systems, which combine multiple artificial intelligence components together with other application components to solve a larger problem, have no known pattern of development and are often approached in a bespoke and ad hoc style. This makes development slower and harder to reuse for future applications. To support the full rapid development cycle of compositional AI applications, we have developed a novel framework called  $(Bee)^*$  (written as a regular expression and pronounced as “beestar”). We illustrate how  $(Bee)^*$  supports building integrated, scalable, and interactive compositional AI applications with a simplified developer experience.

**Index Terms**—rapid development, agile, compositional, artificial intelligence, framework

## I. INTRODUCTION

Delivering applications to users quickly is a cornerstone of agile development methodologies [1], allowing multiple iteration cycles to identify and incrementally solve critical stakeholder problems in response to working applications. However, rapidly developing software requires a developer to focus their effort on what is novel about an application, without distraction from secondary or accidental tasks [2].

Rapidly developing novel systems is difficult today as applications frequently incorporate various types of artificial intelligence (AI) methods such as neural networks, symbolic knowledge representations, reinforcement learning, and others. Compositional AI systems, which combine multiple AI components together with other application components to solve a larger problem, have no known pattern of development and are often approached in a bespoke and ad hoc style. This makes development slower and/or error prone as time needs to be spent on design or coding with no clear architecture, distracting the developer’s attention from critical features.

Tools today support rapid development of AI programs to some extent. For example, Streamlit [3] and Gradio [4] offer a solution to rapidly create frontends by enabling the developer of a Python program (the most common language of AI) to automatically display static data in a web browser and interact with it through a variety of widgets. However, they do not solve several critical challenges in creating compositional AI applications [5], [6]. They offer no support for **integrating** multiple AI components, **scaling** these components on the needed infrastructure, and building dynamically updating applications for rich AI and user **interaction**.

To support the full rapid development cycle of compositional AI applications, we have developed a novel framework called  $(Bee)^*$  (written as a regular expression and pronounced as “beestar”). This framework enables the developer to declaratively build a representation of the entire application (specifying AI components, component and user interactions, visualizations, and desired scaling) in a graph structure that  $(Bee)^*$  interprets in order to automatically create and scale the application. This approach supports rapid application development by simplifying the developer’s responsibility to deciding how to compose AI and GUI elements with declarations, leaving the task of operationally composing the application to  $(Bee)^*$ .

To leverage the graph created by the developer and operate at scale,  $(Bee)^*$  follows an agent-oriented [7], [8] approach to programming, where *agents* are the computational components of the program.  $(Bee)^*$  agents are independent run-times that can take messages to execute their code, coordinate with other agents, and interact with the user with *widgets*, through the graph. As the entire specification of the application is on the graph, widgets and agents can update the application by updating the graph and the  $(Bee)^*$  runtime will dynamically update the application. Further, the graph supports meta-programming for very dynamic behavior, where agents and widgets use it to change source code at runtime.

In Section II, we use a compositional AI application built using  $(Bee)^*$  to show how  $(Bee)^*$  supports **integrated**, **scalable**, and **interactive** compositional AI applications with a simplified developer experience. In Section III, we discuss the feasibility of building with  $(Bee)^*$ . In Section IV we discuss related work. In Section V and VI we discuss future plans and conclude.

## II. $(Bee)^*$ SYSTEM

$(Bee)^*$  uses a graph representation to create and dynamically update applications for compositional AI. Figure 1 shows a GUI (called a Dashboard) of a compositional AI application created with such a graph. Figure 1, ❶, shows a gallery of training data, where each image has been labeled as “bulldozer” (green frame) or not (red frame) by the AI model CLIP [9], or by the user correcting CLIP by clicking on an image to toggle the label. The CLIP agent is run to label the images as “bulldozer” in response to the user typing this

Work licensed under Creative Commons Attribution 4.0 License.  
<https://creativecommons.org/licenses/by/4.0/>

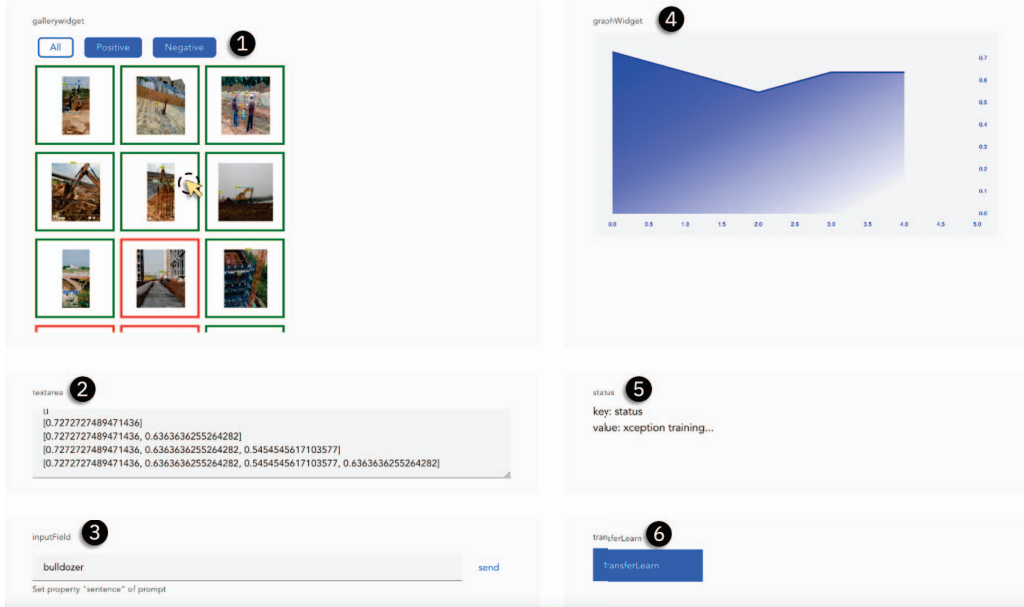


Fig. 1. Dashboard Example.

word into the prompt (Figure 1, ③). Once the user is satisfied with the labels, the user clicks the TransferLearn button (Figure 1, ⑥) to trigger another process that fine-tunes the Xception model [10] on the labeled data so that it can predict bulldozer images (validation accuracy per epoch shown in 1, ④). Status (Figure 1, ⑤) and logging (Figure 1, ②) widgets show updates from the AI components.

To create the application in Figure 1, a graph is specified and executed by the  $(Bee)^*$  architecture (Figure 2). The major components are (1)  $(Bee)^*$  **program**, written by the developer, which declaratively specifies a graph, (2)  $(Bee)^*$  **graph** itself (hosted on some graph database), (3) **agents** that are standalone run-times executing different algorithms/programs and interacting through the graph, and the (4) **Dashboard** that is populated with widgets, specified in the graph, that the user can, in turn, use to view/update the graph.

### A. Graph Representation

We describe a  $(Bee)^*$  graph through an example (Figure 3) used to create the application shown in Figure 1. Nodes in a graph are entities, which specify arbitrary concepts (similar to objects in object oriented programming [11]) with properties and values. For example, in Figure 3 we have an entity named *Training Data* with property *data* and value *link*, where *link* names where the value is stored (in this case we assume a path on disk). This entity will be used by the AI components (i.e., agents) for training and inference. Further,  $(Bee)^*$  provides a primitive typing system, so we know that whatever value *data* is, it is an array.  $(Bee)^*$  knows the entity type of *Training Data* using the edge with name “is a” to the entity named Entity.

1) *Integration with Agents and AgentEntities*: In  $(Bee)^*$ , all processes/algorithms/AI components are agents [7], where an agent is a standalone Python runtime that can take messages (using a remote procedure call server [12]) and map messages to different behaviors. When an agent receives the message “play” it runs its assigned source code, “stop” will stop execution, and “debug” will run assigned source code with the Python debugger [13]. The agent itself is specified as an entity (i.e., AgentEntity) that  $(Bee)^*$  uses to start the agent and as a place to keep and update properties of the agent. The agent also uses its AgentEntity specification to locate its source code to execute on the message “play”. By specifying the source code as a property of the agent, the agent is able to reflect on its own source code as data or other agent’s source code to support meta-programming [14]. Use cases for this include code injection for logging [15], code optimization [16], or genetic programming [17] for functional changes. Further, source code as a property supports the user to debug and change the agent’s source code at runtime with widgets described in Section. II-A2.

AgentEntity inherits from Entity and has the properties *source code* (defining a function), *input* and *output* that specify an agent that  $(Bee)^*$  can launch and run. In Figure 3, there are AgentEntities *CLIPAgent* and *CNNAgent*, where *CLIPAgent* has *source code* to run the CLIP model [9] to label arbitrary images. *CNNAgent*’s source code trains a CNN model (Xception [10]). When an agent executes its source code, it runs by passing the value of its *input* property to the function and maps the return value to its *output* property in its AgentEntity. In this way, the source code, input, and output are “lifted out” of the agent’s runtime and usable and modifiable by any other agents at the graph level, where the agents could be running

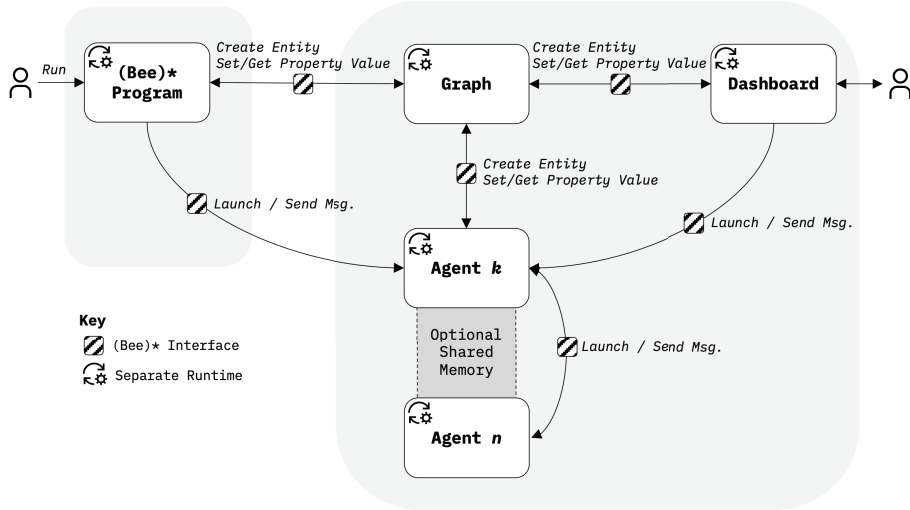


Fig. 2. *(Bee)\** Architecture.

very different source code and living across different servers on a cluster. As such, agents are able to run code in different languages (with operating system calls) but **input and output are independent of their language, supporting integration across languages**. This is helpful in scenarios, for example, where we need to integrate probabilistic programming in Julia [18] with a model in Python.

2) *Display Entities and Input Entities for User Interaction:* *(Bee)\** provides what is called the Dashboard to the user to interact with and view the state of the graph. Widgets on the Dashboard are specified by the developer and, just like agents, widgets are specified as Entities. The developer can specify DisplayEntities (widgets for viewing), InputEntities (widgets for changing properties on Entities), and ButtonEntities to send messages to agents. Custom widgets are possible by extending these types. Once these Entities are declared, the Dashboard finds them on the graph and populates its interface with them. In Figure 3, there is a *TrainDataGallery* entity of type *GalleryEntity* that the Dashboard uses to setup the view of construction data in Figure 1, ①, and customizes its background and border using the *TrainDataGallery*'s properties. A *GalleryEntity* can display any images on the graph by adding a relation specified by an edge (described in next section) from it to the data source, allowing for changes of data sources at runtime and dynamic behavior.

3) *Agent and User Interaction with Watches and Sets:* Any kind of DisplayEntity (e.g., *GalleryEntity*) displays data of interest by specifying an edge to the entity with the property of interest, where the label on the edge is “watches  $\langle$ property name $\rangle$ ”. With this edge, the DisplayEntity is notified to update on any change of that property's value. For example, *TrainDataGallery* displays the training data because it “watches data” of the *TrainingData* Entity. When the data is updated, *(Bee)\** notifies the *TrainDataGallery* to also update. Similarly, an agent is triggered to run its source code if some

property it is interested in changes. In Figure 3, if the property *word* of *Prompt* is changed, then *CLIPAgent* is triggered to run its source code because it is watching that property. When triggered, *CLIPAgent*'s input property is set to the value of *word* and passed as an argument to *CLIPAgent*'s source code.

To support modularity, the output value of an agent can update the value of other properties with the *sets* relationship. In our example, there is an edge from *CLIPAgent* to *Training Data* with label “sets data”, which will set the value of the data property of *Training Data* to the output of *CLIPAgent*. Using this relation has the advantage that it decouples AgentEntities and DisplayEntities. For example, *TrainDataGallery*, will update with new *data* without being coupled to the *CLIPAgent* agent. Further, InputEntities can also set values of properties. In Figure 3, we have an InputEntity, *CLIPInputEntity* (displayed in Figure 1, ③), that sets the value of *word*. Since *CLIPAgent* is watching the value of *word*, it is triggered to run when the user typed in “bulldozer” and not tied to the *CLIPInputEntity* itself.

Continuing with the example in Figure 3, when the value of *data* is set, the user clicks the *TransferLearn* button (Figure 1, ⑥, but not shown in graph example) to send a message to the *CNNAgent* to play its function (specified in its source code) with *data* as *input*. In this case, a CNN model (Xception) used for classification is fine-tuned with *data*, and the weights of the model are stored as type *beestar.tensor* in the *CNNAgent*'s *output* property. A DisplayEntity of type *GraphEntity* (Figure 1, ④, but not shown in graph example) updates with the validation accuracy per epoch during training.

## B. Scale and Efficiency

Agents themselves run in parallel, supporting parallel computation, and can be automatically deployed locally as processes or pods on Kubernetes [19]. Deploying agents on Kubernetes supports the cases where an agent needs a scaled

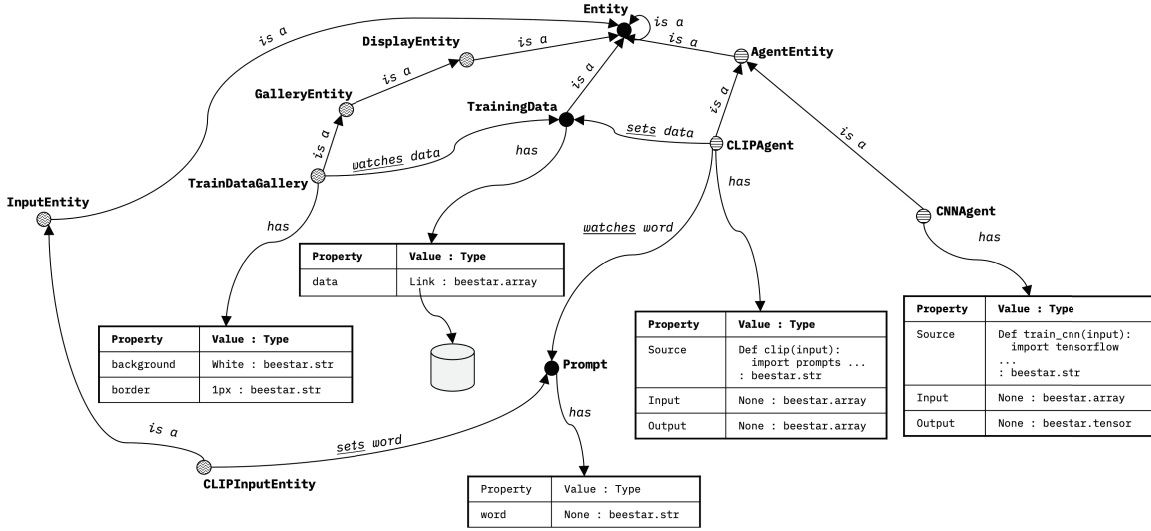


Fig. 3. Part of Graph for Compositional AI Application Example.

```

1 from beestar.agent_functions import clip
2 prompt = Entity("prompt")
3 input = InputEntity(name="CLIPInputEntity")
4 input.sets(prop="word", entities=[prompt])
5 agt = AgentEntity(name="CLIPAgent", func=clip)
6 agt.watch(prop="word", entities=[prompt])

```

Listing 1. Example in  $(Bee)^*$  API

environment and high bandwidth (e.g., GPUs and  $\geq 10$  Gbps network). For Kubernetes,  $(Bee)^*$  uses a predefined base container to create an agent pod on Kubernetes, installs the needed requirements from the agent’s requirements property, and launches the agent inside the container (mapping ports as needed at the cluster level). From the developer’s perspective, they simply change a parameter to deploy locally or on Kubernetes and specify requirements in the agent’s property, making it easy for the developer to scale their application.

### C. $(Bee)^*$ Interface and Developer Experience

While the graph can grow in complexity, all the specification and update logic is handled through  $(Bee)^*$  Interface library to support a simpler developer experience and enforce the watches and sets relationships. For declarative specification, the developer declares they would like an agent, widgets, entities, and update rules by calling methods in the  $(Bee)^*$  Interface. Listing 1 shows the sequence of calls to create a prompt Entity, an InputEntity that sets the *prompt’s* word after the user types in their keyword, and a *CLIPAgent* that watches the word or prompt. With six lines, the  $(Bee)^*$  runtime is able to create a web application that takes a word from an input textbox and passes the word to a large foundational model (CLIP) so that it can run inference.

The  $(Bee)^*$  Interface library also handles update rules. When any property value is being updated, it has to happen

through the  $(Bee)^*$  Interface (as shown with the boxes on the edges in Figure 2). On property change, the  $(Bee)^*$  Interface queries the graph for sets/watches relationships and, in turn, updates values and notifies watchers as needed. The  $(Bee)^*$  Interface library need not be centralized as long as the update rules are consistent across all instances.

### III. FEASIBILITY OF DEVELOPING WITH $(Bee)^*$

We looked at how efficient the application in Figure 1 is, to understand overhead in  $(Bee)^*$ , and also explored a breadth of game applications that can quickly be played by agents. When running the  $(Bee)^*$  application in Figure 1 on a machine with one P100 GPU on Kubernetes, we found that the *CLIPAgent* was able to label 4000 construction images in 52.1s and the *CNNModel* agent trained in 111.7 seconds. Since labeling is a labor intensive tasks, the speed of the agents speaks to the feasibility of  $(Bee)^*$  to create programs that are fast enough to be useful.

We observed how feasible it might be to build a variety of agents working together to play arbitrary games in Figure 4. In Figure 4, ①, the DOS [25] game called The Incredible Machine [20] is being played by a variety of agents. Figure 4, ③ is a code editor widget for the *VNCagent* that takes screen shots of the game (Figure 4, ④), and Figure 4, ⑤ is a code editor widget for a *ObjectDetectorAgent* that identifies objects from *VNCagent’s* screen shot. The objects detected are set as a value for the *Game* Entity, which the gallery widget (Figure 4, ②) watches and its images update as the objects update. From this  $(Bee)^*$  program, we were able to use the same agents and widgets to play Monkey Island [21], Doom [22], Donkey [23], and Manic Mansion [24], by only changing the URL of where the game is being hosted in the graph. Further, the editors for the agents let the programmer tweak the code in real time for each game, showing the value meta-programming.

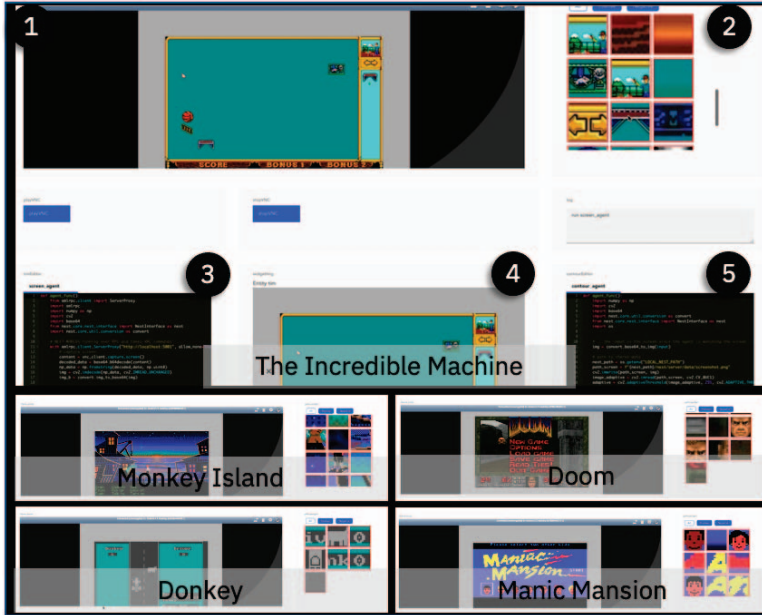


Fig. 4. Same  $(Bee)^*$  program different games [20]–[24].

#### IV. RELATED WORK

Borowski, et al. present the Varv [26] declarative programming language that supports building a website live in the browser, where components live in the DOM [27] of the browser. Similarly, Perez De Rosso et al. present a declarative framework for websites with components called concepts [28]. While  $(Bee)^*$  is declarative and supports dynamic applications, it differs in that it can declaratively integrate AI components, where the AI could be in other languages, frameworks, and infrastructure that are outside the browser, and, additionally, it can integrate agents with GUI elements in a browser. Together, these features provide a framework for integration, interaction, and scale of AI components not previously supported.

In the AI literature, using a graph to integrate agents has been demonstrated by Goertzel, et al. with the OpenCog framework [29], where the focus is on collaboration among agents. Our work addresses using a graph to also integrate application components (e.g., GUI components).

Low code frameworks (e.g., Node-Red [30], ConveyorAi [31], and Patterns [32]) use a directed graph to specify the execution and information flow of services. In contrast, we present a graph that is a shared knowledge representation, using a blackboard architecture, that is used and reflected on by agents and users for interaction and collaboration, but, further, includes components to specify a rich user interface.

#### V. FUTURE PLANS

Future work includes both an in depth evaluation of  $(Bee)^*$  to rapidly prototype compositional AI applications and also demonstrations of using the framework to further illustrate the power of the approach of agents collaborating over the graph.

For evaluation, we will measure the speed developers can create AI applications with  $(Bee)^*$  versus popular approaches today. For demonstration, we will show agents searching the graph for other agents to “outsource” work to them, show gradients passed across agents so they learn together, and show public graphs to share agents and widgets across developers.

#### VI. CONCLUSIONS

$(Bee)^*$  is a novel rapid development framework for compositional AI applications. We demonstrated how a developer can declaratively specify agents, widgets, and their relationships for rapidly integrating and scaling dynamic compositional AI applications. With the sets and watches relationships between agent and widget entities,  $(Bee)^*$  provides a dynamic agent and user **interaction**. For **integration** across different kinds of agents,  $(Bee)^*$  integrates agents and widgets at properties in the graph and not inside the agent.  $(Bee)^*$  supports **scale** and parallelism by running agents in parallel and on Kubernetes.

We looked at examples to explore the feasibility of building compositional AI with  $(Bee)^*$ . We showed CLIP and CNN models collaborating as agents to label thousands of images and train a CNN model in minutes. Further, we observed the modularity of  $(Bee)^*$  by showing game playing agents cooperating across five different DOS games, where only the URL of the game changed across each of the applications.

#### VII. ACKNOWLEDGEMENTS

We would like to thank Grady Booch, John Cohn, Adriana Meza Soria, and Ryan Anderson for their input. We also thank the MIT-IBM Watson AI Lab for supporting this work.

## REFERENCES

- [1] J. Shore and S. Warden, *The art of agile development*, 1st ed. O'Reilly.
- [2] Brooks, "No silver bullet essence and accidents of software engineering," vol. 20, no. 4, pp. 10–19, conference Name: Computer.
- [3] Streamlit • the fastest way to build and share data apps. [Online]. Available: <https://streamlit.io/>
- [4] G. Team. Gradio. [Online]. Available: <https://gradio.app/>
- [5] Is it possible to update the UI dynamically ? · discussion #404 · gradio-app/gradio. [Online]. Available: <https://github.com/gradio-app/gradio/discussions/404>
- [6] How to prevent the reloading of the whole page when i let the user to perform an action - using streamlit. [Online]. Available: <https://discuss.streamlit.io/t/how-to-prevent-the-reloading-of-the-whole-page-when-i-let-the-user-to-perform-an-action/10800/3>
- [7] N. R. Jennings, "Agent-oriented software engineering," in *Multi-Agent System Engineering*, ser. Lecture Notes in Computer Science, F. J. Garlajo and M. Boman, Eds. Springer, pp. 1–7.
- [8] B. Hayes-Roth, "A blackboard architecture for control," vol. 26, no. 3, pp. 251–321. [Online]. Available: [https://doi.org/10.1016/0004-3702\(85\)90063-3](https://doi.org/10.1016/0004-3702(85)90063-3)
- [9] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *Proceedings of the 38th International Conference on Machine Learning*. PMLR, pp. 8748–8763, ISSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v139/radford21a.html>
- [10] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1800–1807, ISSN: 1063-6919.
- [11] G. Booch, *Object Oriented Design with Applications*. Benjamin/Cummings Publishing Company, google-Books-ID: w5VQAAAAMAAJ.
- [12] J. Bloomer, *Power programming with RPC*. O'Reilly & Associates, Inc.
- [13] pdb — the python debugger — python 3.10.7 documentation. [Online]. Available: <https://docs.python.org/3/library/pdb.html>
- [14] K. Czarnecki, K. Østerbye, and M. Völter, "Generative programming," in *Object-Oriented Technology, ECOOP 2002 Workshops and Posters, Málaga, Spain, June 10-14, 2002, Proceedings*, ser. Lecture Notes in Computer Science, J. H. Núñez and A. M. D. Moreira, Eds., vol. 2548. Springer, pp. 15–29. [Online]. Available: [https://doi.org/10.1007/3-540-36208-8\\_2](https://doi.org/10.1007/3-540-36208-8_2)
- [15] L. Bergmans and C. V. Lopes, "Aspect-oriented programming," in *Object-Oriented Technology ECOOP'99 Workshop Reader*, ser. Lecture Notes in Computer Science, S. Demeyer and A. Moreira, Eds. Springer, pp. 288–313.
- [16] A. V. Aho, M. S. Lam, R. Sethi, and J. D. Ullman, *Compilers: Principles, Techniques, and Tools (2nd Edition)*. Addison-Wesley Longman Publishing Co., Inc.
- [17] W. Banzhaf, F. D. Francone, R. E. Keller, and P. Nordin, *Genetic programming: an introduction: on the automatic evolution of computer programs and its applications*. Morgan Kaufmann Publishers Inc.
- [18] The julia programming language. [Online]. Available: <https://julialang.org/>
- [19] Production-grade container orchestration. [Online]. Available: <https://kubernetes.io/>
- [20] The incredible machine. [Online]. Available: <https://www.myabandonware.com/game/the-incredible-machine-1mg/play-1mg>
- [21] Download the secret of monkey island | DOS games archive. [Online]. Available: <https://www.dosgamesarchive.com/download/the-secret-of-monkey-island>
- [22] Doom. [Online]. Available: <https://www.myabandonware.com/game/doom-1nd>
- [23] Donkey. [Online]. Available: <https://www.myabandonware.com/game/donkey-1r0>
- [24] Download maniac mansion | DOS games archive. [Online]. Available: <https://www.dosgamesarchive.com/download/maniac-mansion>
- [25] DOSBox, an x86 emulator with DOS. [Online]. Available: <https://www.dosbox.com/>
- [26] M. Borowski, L. Murray, R. Bagge, J. B. Kristensen, A. Satyanarayan, and C. N. Klokose, "Varv: Reprogrammable interactive software as a declarative data structure," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI '22. Association for Computing Machinery, pp. 1–20. [Online]. Available: <https://doi.org/10.1145/3491102.3502064>
- [27] J. Keith and J. Sambells, *DOM Scripting: Web Design with JavaScript and the Document Object Model*, 2nd ed. Apress.
- [28] S. Perez De Rosso, D. Jackson, M. Archie, C. Lao, and B. A. McNamara III, "Declarative assembly of web applications from predefined concepts," in *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, ser. Onward! 2019. Association for Computing Machinery, pp. 79–93. [Online]. Available: <https://doi.org/10.1145/3359591.3359728>
- [29] B. Goertzel, C. Pennachin, and N. Geisweiller, "The OpenCog Framework," *Engineering General Intelligence, Part 2*, pp. 3–29, 2014, publisher: Atlantis Press, Paris.
- [30] Node-RED. [Online]. Available: <https://nodered.org>
- [31] Conveyor AI. [Online]. Available: <https://conveyorai.com/conveyorai.com>
- [32] Patterns | Build next-gen AI systems | Patterns. [Online]. Available: <https://www.patterns.app>