

Marvin-AI

Marvin-AI

Abstract

Marvin-AI is an open-source artificial intelligence (AI) platform that helps data scientists, prototype and productionalize complex solutions with a scalable, low-latency, language-agnostic, and standardized architecture while simplifies the process of exploration and modeling.

Proposal

Marvin helps non-experienced developers create industry-grade AI applications. It has three core components: a development environment to be used during data exploration and hypothesis validation (Toolbox), a library which should be extended to create Marvin engines, and a Scala application server which interprets engines (Engine Executor). A basic premise of Marvin is that it should be language-agnostic, able to interpret engines implemented in different programming languages.

Background

The Marvin AI project was initiated as an internal project at B2W Digital (Brazil), the largest e-commerce company in Latin America. Nowadays, it is used by all data scientists within the B2W team. Oftentimes, data scientists don't have an extensive background in software engineering, yet are in charge of creating AI applications that need to scale to high throughput and provide millisecond-level response times. At B2W, Marvin AI plays an important role in this process, abstracting advanced software engineering procedures, allowing data scientists to focus on their knowledge domain.

Rationale

With recent advances in computer architecture and a corresponding increase in the amount of data generated by always-connected devices, AI algorithms offer a solution to problems that have long troubled modern corporations. Since AI developers come from various fields, such as statistics, physics, and math, there exists a strong need for platforms which enable them to move from prototypes to enterprise applications. Although some tools claim to offer this service, in reality, there is no reliable open-source solution.

Initial Goals

The initial goals will most likely be to merge the existing codebase into a single repository, migrate it to Apache, and then integrate with the Apache development process. Furthermore, we plan for incremental development and releases, as per Apache guidelines.

Current Status

Meritocracy

Marvin already works under principles of meritocracy. Today, Marvin already has some contributors that are part of other institutions. Although there is no formal process defined to become a committer, contributors that make major changes/improvements to the platform are naturally granted write access to the repository.

Community

Acceptance into the Apache foundation would substantially boost both Marvin's user and developer communities. The current community includes a few experienced developers that have either academic or professional experience with AI. The community is largely comprised of data scientists working at B2W and other companies such as Cloudera, MIT, Qume Labs, Laguro.com, and CBYK. Also, there is a meetup group of hundreds of users who meet regularly to exchange ideas about Marvin and, more generally, AI.

Reference to the group: <https://www.meetup.com/marvin-ai/members/>

Core Developers

The core developers for Marvin are listed in the contributor's list and initial PPMC below. These lists include B2W employees, MIT students, UFSCAR researchers, independent contributors, and some employees of other companies like Cloudera, Qume Labs, Laguro.com, and CBYK.

Alignment

The initial committers strongly believe that by being part of the Apache Software Foundation, Marvin AI will be part of a comprehensive suite for AI applications that can process big data and enable enterprises to extract value from their data lakes. Also, we hope that by integrating with other Apache projects such as Apache Spark, Apache Hadoop; that this will foster additional collaboration between these projects furthering the already existing integration points and expanding the community of contributors.

Known Risks

Orphaned products

Given the current maturity of Marvin and how well it has been received at technical conferences, the risk of the project being abandoned is minimal. AI is not academia-exclusive anymore, and as enterprises start to add data-science pipelines to their applications, demand for Marvin will only increase.

Inexperience with Open Source

Marvin AI has been an open-source project since October 2017. The project was started in a company where open-source culture is foundational. B2W Digital runs the largest e-commerce in Latin America on top of open-source projects.

Reliance on Salaried Developers

Marvin AI receives substantial efforts from salaried developers – a few of which were hired by companies to work exclusively for the project – but the majority devote "after-hours" or spare time to this project. Some developers are graduate students that contribute in their free time at school.

Relationships with Other Apache Products

Marvin integrates with several Apache products, such as Hadoop (HDFS) and Spark. Marvin shares some similar features with PredictionIO, specifically the model application server and a design pattern that was inspired by the DASE. Despite these similarities, Marvin is catered towards a different clientele (data scientists), and for that reason, it includes many critical features that are not provided by PredictionIO.

An Excessive Fascination with the Apache Brand

While the ASF brand will undoubtedly help Marvin become a successful project, Marvin is already gaining traction at companies around the globe.

Documentation

<http://www.marvin-ai.org>

Initial Source

The current codebase is available at <http://github.com/marvin-ai>. This is practically the same code that will be migrating to the Apache Foundation, the notable difference being that the multiple repositories will be merged into a single repository (if necessary).

These are the main repositories and a very simplified explanation about each one:

Main repositories

- [marvin-ai/marvin-python-toolbox](#) - Data Science toolbox that helps in the creation of new ML engines
- [marvin-ai/marvin-engine-executor](#) - Component responsible for interpreting, serving and managing Marvin engines
- [marvin-ai/marvin-public-engines](#) - Marvin engine examples to help new Marvin users to build engines
- [marvin-ai/marvin-platform-book](#) - Documentation in [GitHub](#) book site format

Secondary repositories (Experimental and Initial)

- [marvin-ai/marvin-vagrant-dev](#) - Development environment that uses [VirtualBox](#) and vagrant to non mac and Linux users;
- [marvin-ai/marvin-paper](#) - Source code (latex format) of the first Marvin paper published in PAPIS.io conference in Boston.
- [marvin-ai/marvin-cluster-admin](#) - Admin module responsible to manage Marvin cluster;
- [marvin-ai/marvin-automl](#) - AutoML module responsible to help data scientist to build machine learning models with a very simple visual interface;

External Dependencies

It is very likely that all our dependencies are using either the Apache or MIT license. Upon acceptance to the incubator, we would begin a thorough analysis of all transitive dependencies to verify this fact and introduce license checking into the build and release process.

Required Resources

Mailing lists

- private@marvin.incubator.apache.org (with moderated subscriptions)
- dev@marvin.incubator.apache.org
- commits@marvin.incubator.apache.org

Git Repositories

- <https://git-wip-us.apache.org/repos/asf/incubator-marvin.git>

Issue Tracking

- JIRA (MARVIN)

Initial Committers

- Lucas Bonatto Miguel <lucasbonatto@gmail.com> - Qume Labs (California - USA)
- Daniel Takabayashi <daniel.takabayashi@gmail.com> - B2W Digital (São Paulo - BR) / Laguro.com (California - USA)
- Bruno Piraja <bruno.piraja@b2wdigital.com> - B2W Digital (São Paulo - BR)
- Zhang Yifei <zhang.yifei@b2wdigital.com> - B2W Digital (São Paulo - BR)
- Harrison Wang <hwang123@mit.edu> - MIT (USA)
- Brody West <brodyw@mit.edu> - MIT (USA)
- Rafael Novello <rafael.novello@b2wdigital.com> - B2W Digital (São Paulo - BR)
- Willian Leite <willian.leite@cbyk.com.br> - CBYK (São Paulo - BR)
- Danilo Nunes <nunesdanilo@gmail.com> - Qume Labs (California - USA)
- Alan Silva <alan.silva@cloudera.com> Cloudera (USA)
- Jeremy Elster <jeremy.elster@b2wdigital.com> - B2W Digital (São Paulo - BR)

Sponsors

Champion

- Luciano Resende - (lresende)

Nominated Mentors

- Luciano Resende - (lresende)
- Jim Jagielski - (jim)
- William Colen - (colen)

Sponsoring Entity

We would like to propose the Apache Incubator to sponsor this project.