Department of Computer Science and Engineering, University of Nevada – Reno

Gigafactory Systems Machine Learning Project

Team 15: Adam Cassell, Braeden Richards, Ashlee Ladouceur

Instructors: Sergiu Dascalu, Devrin Lee

External advisor(s): Dr. Emily Hand (UNR), Gavin Hall (Tesla)

8 February 2019

Table of Contents

Abstract

Team 15 is creating a machine learning system that provides real-time and offline analysis of Tesla's Gigafactory systems. The pipeline of the system design includes building an interface for users to connect to Tesla's mySQL database containing both historical and current data of sensors in the factory, building models, training and testing those models, and building a visual analysis of the data received from the training and testing. This particular document provides a detailed project concept and an outline of the project management and the associated responsibilities. Team 15 has begun development of the user interface and has started to preprocess sensor data for model building.

Project Description

Team 15 will provide the Tesla Gigafactory with a full stack application that will be aimed at two categories of user: machine learning engineers and other engineers. GigaML will provide an invaluable service to the Tesla Gigafactory by providing tools to increase efficiency of environmental and factory systems as well as reduce workload for current factory modeling engineers. The application will provide a service for the creation and training of machine learning models as well as provide services for the viewing and analysis of model/system data and preprocessing of system data.

The application will have two distinct dashboards, one for administrators and one for normal users. Administrators will be able to create new models or modify current models as well as use the application's base features. Normal users will only have access to the base features, such as displaying system/model data from the servers, graphing past and present data, and using stored models on system data. Having a split in user groups maintains model integrity and ensures model efficiency by only allowing qualified administrators, such as machine learning engineers, to train new models or remove current models. This also ensures that the application will have a reasonable number of accurate models instead of many partially-trained or incomplete models made by the average user.

GigaML will have three distinct pieces: frontend, backend, and the MySQL server. The frontend will be hosted on a Windows server and will be created using a combination of the .NET framework, JavaScript, and HTML/CSS. The front end will be intuitive and easy to use for the average user and will act as the trigger for the rest of the application. The UI will be able to obtain data for presentation from the MySQL server and will call the backend Python scripts to complete any preprocessing, data manipulation, or machine learning that the user requests. The backend will be hosted on a Linux server and will be created using Python3.6, TensorFlow with Keras, and NumPy. The users of the application will not have direct access to the backend but will instead communicate with it using the frontend UI. The backend will be able to obtain, modify, and send data to the MySQL server and send and receive data from the frontend. Administrators will have access to the backend for maintenance and for the addition or removal of scripts during updates to the software.

To obtain the goal of secured data and a secure application, GigaML will be using the admin/base user system to give different security restrictions depending on the user. The application will also only be used within Tesla's systems and will only be accessible from within their network, so all users will have passed security checks within Tesla. To reach the goal of dependability, the software will be designed with checks and tests to minimize the risk of crashing and will include flags and locks to ensure data will not be modified during reading, training, etc. Reliability will also be ensured with Keras and NumPy, which are large and heavily supported libraries that are efficient with resources and reliable.

Significance

GigaML is a highly valuable educational and professional opportunity for Team 15. Team 15 will obtain full stack development experience and opportunities for future learning through research and the creation of the software. For Tesla, the project will provide a basis to build models that could increase their efficiency with resources and time and provide an opportunity for decreasing costs across the Gigafactory.

For Team 15, the opportunity for professional growth during this project is substantial. Team 15 will be able to put their skills that have been picked up through their educational career to the test by building a full stack application. Furthermore, by being advised by Gavin Hall, a machine learning engineer at Tesla, and Emily Hand, a machine learning Professor at UNR, Team 15 has great access to learning resources and mentorship. Team 15 will also be able to gain important professional connections through their close work with Tesla as well as with the machine learning community at UNR. Once the project is completed, each member of Team 15 will have the ability to discuss the project with potential employers as an invaluable experience that will put them ahead of the average software engineer.

An application like GigaML has never been used by Gigafactory engineers for system design. Currently, most systems are controlled by modeling engineers and the systems that are controlled by machine learning models are custom designed by Gavin Hall. GigaML will allow for models to be created more efficiently and provide a way for other engineers to use the models. Since system specific well-trained machine learning models are often more performant at making decisions than humans, this has the opportunity to provide more accurate results for the Gigafactory – saving money, time, and energy. GigaML is being designed solely for use by Tesla, so there is no other market potential for the application

When GigaML is completed there will still be lots of room for future software expansion and feature addition. At release, GigaML is expected to only support neural networks as a model type for creation and analysis. This is because many of the systems used at Gigafactory have demonstrated successful modeling predictions using neural networks in Gavin's prior work. In the future, however, updates to add more model types, such as Naïve Bayes trees, Random Forests, or other decision trees could be added to grow the application. The application at launch is only guaranteed to have one trained model already in it, and that will be trained on a specific system by Team 15. Therefore, more trained models should be added for base users to use as the application matures past release.

Legal and Ethical Aspects

GigaML will strictly comply with all ethical and legal expectations. Development will satisfy the ACM code of ethics PRODUCT clause in several ways. First, Team 15 plans to communicate our development schedule on a regular basis, including any components that could potentially become ahead of or behind schedule (3.01). The team will negotiate any timeline changes with Tesla and make sure all parties involved understand the benefits and risks in doing so. GigaML will be designed at every step of the way with the highest quality and programming practices in mind. Team 15 will develop the software in a way that is cost-efficient, using open-source and freely available software whenever possible. Clear specification documents have been and will be created outlining product requirements and outlines (3.08). Team 15 will also implement automated tests on all frontend and backend components of the software (3.10). The team will consistently document all progress and challenges (3.11). All data used in the software will only be derived with explicit permission from Tesla and using the appropriate and approved channels (3.13). Any data that seems inaccurate or not representative of correct factory system data will be flagged and handled appropriately so as to not adversely affect engineering decisions (3.14).

Legally, Team 15 has an understood agreement when it comes to the development and use of the software. The software will not be distributed outside of the development team and Tesla. For university presentations, and anywhere else where there is public visibility, only generalized content and details will be discussed as approved by Tesla. Any factory data or other confidential information, alongside the final software, is the property of Tesla alone. Equipment provided to team 15 is likewise the property of Tesla and will be returned to the company at the completion of the project.

Changes and Progress

The team has made significant developmental progress since delivering the prototype. This includes further refining the User Interface (UI) in accordance to the external advisor from Tesla's feedback during the user feedback session held on 12/5/2018. In addition, the team has received a sample of sensor data provided from the Gigafactory, allowing for the team to begin writing the Python backend and preprocess the data for model building. Finally, the team has fully established the main project components needed to reach a deliverable by May. This was important in order to have a guideline for the team to follow in order to focus on the system priorities. The responsibilities of each component have been split between team members to focus on in order to ensure a productive development.

A major change made to the project since the concept in early October 2018 is the difference in user experience. The concept originally had separation between a user and admin, and now the project has included different user and admin capabilities. Due to this, the MySQL server has included an additional database to hold user information provided from a Tesla Gigafactory employee roster. This change was necessary in order to not allow users to build models, but only allow admins who are trained in the system. Users will only be able to view the data from the models the admins build. The team does not anticipate any further changes thanks to the heavy positive feedback from Gavin Hall, the team's external advisor from Tesla, and in establishing a detailed project outline for the team to follow throughout development.

In summary, Team 15 has achieved the following major developments and accomplishments in creating the GigaML system:

- Received Gigafactory sensor data from Tesla
- Began development on the Python backend, including preprocessing data received from Tesla so TensorFlow can begin building models to be trained and tested
- Refined the UI layout according to specifications provided from Gavin Hall
- Began planning database layout to hold sensor data in the most effective way
- Set a weekly four-hour development meeting to begin building the system and ensure timely delivery
- Distributed responsibilities of project components among team members

Project Responsibilities

The main components of the delivered product can be split into three main systems: the User Interface (UI) frontend, the Python backend, and the MySQL server. These three systems each contain components that Team 15 will split up evenly in management to ensure high-quality deliverables. It is understood that when a team member is in charge of a component, they are responsible for its concept, specification and design, implementation, and testing. It is important to note that the team will be utilizing pair programming in implementing the systems so collaboration will be used in order to increase knowledge share and ensure consistent coding style.

The UI frontend will be what the users will interact with in order to interact with the system functionalities. Table 1 describes the main components of the UI and the associated team member in charge.

Component	Description	Team Member
Permissions	Reads from the MySQL server the user data and the associated names and permissions	Adam
User experience	Users will have limited access to the functionalities of the system	Adam
Admin experience	Admins will have access to all functionalities of the system	Adam
Initiate Python script	The UI will be accurately tied to the Python backend	Braeden
Data output	Reads from the MySQL server the model information to display to both users and admins	Adam

Table 1: Detailed User Interface subcomponents and their associated team member.

The Python backend will be the largest implementation in developing the system. The backend will contain the main functionalities of the overall system. Table 2 describes the main components of the Python backend and the associated team member in charge.

Component	Description	Team Member
Data I/O	Reads from the MySQL server the sensor data, outputs data back into the MySQL server the model information	Braeden
Preprocessing	Preprocesses the sensor data for model building	Braeden
TensorFlow	Uses preprocessed data to build models	Adam
Model Training	Trains the models built by TensorFlow	Braeden
Model Testing	Tests the accuracy of modelss	Braeden

Table 2: Detailed Python backend subcomponents and their associated team member.

The MySQL server will contain all of the data used by both the UI and the Python backend. On the one server, there will be two databases. The first database will contain the sensor data and the model data. The second database will contain the names and allowances of users. Table 3 describes the main components of the MySQL server and the associated team member in charge.

Component	Description	Team Member
Sensor data	Stores data received from Tesla's Ignition about sensors around the Gigafactory	Ash
Model data	Stores data regarding model architecture received from the backend about models	Ash
User/Admin data	Stores Tesla employee data of names and permissions	Ash

Table 3: Detailed MySQL server subcomponents and their associated team member.

In addition to the detailed subsystems, Team 15 will put high priority on the following components to ensure successful development after the team moves on from the project. This will allow for Tesla to continue any necessary development with ease. Table 4 describes the main system developmental components and the associated team member in charge.

Component	Description	Team Member
Testing	Writing user experiences based off use cases supplied by Tesla to test system functionality during development	Ash
Documentation	Ensuring all code is heavily commented and establishing a base code style for all team members to follow	Ash

Table 4: Detailed overall components and their associated team member.

Project Monitoring and Risks

<u>Risk Register</u>

Risk ID	Risk	Likelihood (1-10)	Impact (1-10)	Date Raised	Mitigation Strategies
1	Losing development resources (i.e. loss of team member)	1	5	1-Feb-19	Ensure all members of development team are familiar with all components of the software to minimize knowledge silos
2	Limitations with Tesla-provided resources and assistance, including advisement time	3	4	1-Feb-19	Plan meetings with advisors well in advance and with key goals for maximum meeting efficiency
3	Limitations on consistent access to factory system data	4	5	1-Feb-19	Design the system so that it can be tested against example subsets of data in the event direct network access is lost
4	Linux Server Delays - Backend development is dependent on a dedicated linux server being commissioned onsite without delay	2	5	1-Feb-19	Coordinate with advisors and Tesla IT to ensure any access limitations to servers are sorted out early in development
5	Limitations in access to Tesla's GPU cluster for model training	2	4	1-Feb-19	Work with advisors to learn more about cluster availability and any restrictions to avoid interfering with other projects
6	Requirements changes and feature suggestion after initial demos could add significant development time	2	5	1-Feb-19	Create and maintain clear specification documents and communicate progress often. Keep expectations realistic.
7	Limitations with testing software off-site. It should be possible to test progress when not physically on the Gigafactory network	5	5	1-Feb-19	Work with advisors and Tesla IT to secure remote vpn access to allow for easier testing of software off-site
8	Noisy data from factory system sensors	6	6	1-Feb-19	Focus dedicated development time on analyzing the data itself and on writing effective pre-proccessing functions.

Figure 1: Risk register for GigaML project.

Team Overview

Team 15

Adam Cassell is a senior computer science student at the University of Nevada, Reno. He is enrolled in CS 426 for senior projects. Adam interned with Tesla for nearly two years while working towards his undergraduate education. He has heavy industry experience with web-application development using JavaScript, aspx.NET, and MySQL. Adam has also taken artificial intelligence and machine learning courses and contributed to undergraduate research on neural network classification. Adam is experienced with Python and TensorFlow which will be critical for this project. Adam will be in charge of web-application development and focus on the front-end interface. Like all team members, he will also assist and maintain familiarity with the other project components.

Ash Ladouceur is a senior at the University of Nevada, Reno enrolled in the CS426 course for senior projects. She is employed at a local software company that used MySQL as their primary programming language. In addition, she has taken a Data Mining and Artificial Intelligence class that focused on data preprocessing and model training and testing, respectively. Ash will be in charge of the MySQL server and the two databases that contain all data used by both the frontend and backend. She will also be in charge of maintaining documentation and testing of the code produced.

Braeden Richards is a senior at the University of Nevada, Reno enrolled in the CS426 course for senior projects. He has experience with Python in both backend and frontend applications. He has taken a course in Data Mining which focused on preprocessing of data and has experience working with large amounts of data in Python. Braeden will be in charge of the backend of this project, ensuring proper practices and reliability when the group creates it.

Advisors

Gavin Hall is in charge of AI systems development at the Gigafactory for Tesla. HIs research projects focus on the application of machine learning algorithms to computer vision, optimal control, and predictive analytics of infrastructure and manufacturing systems. Gavin is team 15's primary point of contact at Tesla and responsible for providing feedback on the end product. He will also be heavily involved in guiding software development progress.

Emily Hand is an assistant professor in the Department of Computer Science and Engineering at the University of Nevada, Reno. Her research focuses on the intersection of computer vision and machine learning and focuses on understanding human perception. For this project, Dr. Hand is providing team 15 guidance on neural network implementation. She is responsible for being an academic resource on neural networks

Team Contributions

Team 15 worked together to create the Revised Concept and Project Management document with the understanding of group effort. A two-hour long meeting was held between the team members prior to creating the document, so all members understood each section goal and outline project development. Table 5 is a representation of the time worked on each section:

Team Member	Sections	Time Worked
Adam Cassell	Legal and Ethical Responsibilities, Project Monitoring and Risks, Team Overview	5.5 hours
Ashlee Ladouceur	Abstract, Changes and Progress, Progress Responsibility, Team Overview, Team Contributions	5 hours
Braeden Richards	Project Description, Significance, Team Overview, References	5 hours

Table 5: Detailed breakdown of time per section by each member of Team 15 for the Revised Concept and Project Management document.

References

"Deep AI." Website. Retrieved October 27, 2018 from https://www.deepai.org.

"Colaboratory." Website. Retrieved October 27, 2018 from https://colab.research.google.com.

"Towards Data Science." Website. Retrieved November 12, 2018 from https://towardsdatascience.com/gradient-descent-in-a-nutshell-eaf8c18212f0.

"Introduction to Hopfield Neural Networks." Website. Retrieved November 12, 2018 from https://www.doc.ic.ac.uk/~sd4215/hopfield.html.