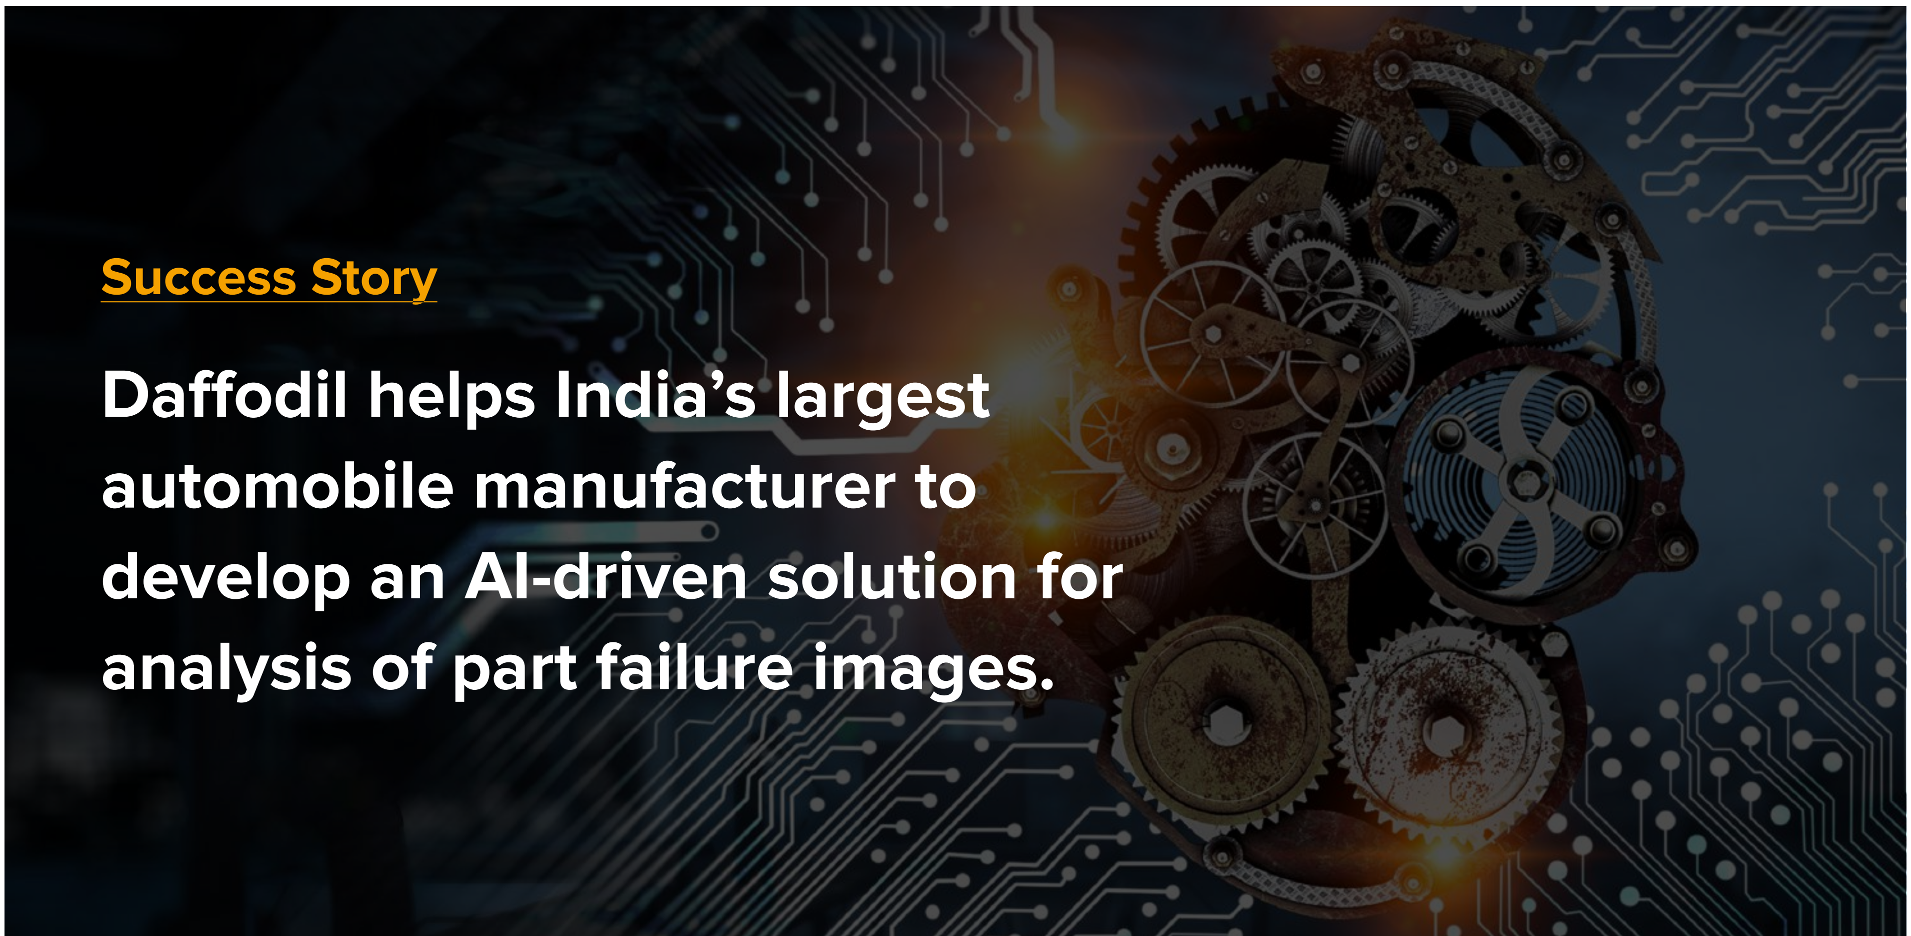

Success Story

Daffodil helps India’s largest automobile manufacturer to develop an AI-driven solution for analysis of part failure images.

Customer	Country	Industry
India’s largest automobile manufacturer	India	Automobile

About The Client

The client is India's largest automobile manufacturer. It is credited with having ushered in the automobile revolution in the country. The company is engaged in the business of manufacturing and sale of passenger vehicles in India. The client has a vast portfolio of 16 car models with over 150 variants. Other activities include facilitation of pre-owned car sales fleet management, car financing.

Technology Stack

React	python

Business Situation

The client, a premier automobile manufacturer, faced significant challenges with its complex and labor-intensive motor part failure and fracture analysis process. The existing system, heavily reliant on expert evaluations, resulted in time-consuming, inconsistent, and inefficient outcomes.

Additionally, the company struggled to efficiently respond to sudden or emergency assessment requirements. The categorization of fractures, whether brittle or ductile, and determining the fracture propagation direction, presented considerable difficulties.

Recognizing these challenges, the client engaged Daffodil Software with a clear objective: to leverage artificial intelligence and image analytics to revolutionize their failure and fracture analysis process. The aim was to develop a solution that would not only expedite the analysis, detection, and reporting of failures but also automate and enhance the accuracy of these analyses.

Current Situation:

- The current process at the client's organization for analyzing fractures in motor parts was meticulous and time-consuming.
- The process heavily relied on the expertise of professionals.
- The existing system posed difficulties in accurately categorizing fractures (as brittle or ductile) and determining the direction of fracture propagation.
- It was not capable of efficiently handling ad-hoc or emergency assessments

Hurdles on the way

The development process presented a multitude of hurdles that demanded strategic navigation and innovative problem-solving. The following elaborates on the specific challenges encountered during the project, emphasizing the complexity of each issue faced by the Daffodil Software team.

Working with a Limited Training Dataset

The team faced a significant challenge related to the volume of data available for preparing machine learning data model. Specifically, the client provided us with a total of 701 images. Additionally, our dataset, consisting of very few images per category, posed difficulties in obtaining meaningful and accurate results.

This was largely due to the fact that machine learning models typically require substantial amounts of data to effectively learn and make accurate predictions. The limited dataset size meant that the models could not learn as effectively as they would with a larger dataset.

To add to the complexity, the team had to classify the data based on several grades, magnification levels, material, and fracture types.

Choosing The Right AI Models

The Daffodil team embarked on a rigorous exploration of 100s of AI models. This extensive experimentation was aimed at identifying the optimal AI model that would best serve the project's unique needs. Some of the models the team tested included U-Net, Deeplabv3, Fast FCN, SegNet, and PSPNet. The process was rigorous and required a careful analysis of each model's performance under these specific conditions.

Annotation hurdles

- **Challenges with annotation tool:** Team Daffodil encountered a series of challenges concerning the selection and usage of an appropriate annotation tool. After careful consideration, the team decided to opt for Amazon Web Services' Ground Truth service for the task of data annotation. This selection provided the team with an opportunity to broaden their skill set and gain hands-on experience with a new and innovative tool.

- **Challenges with Annotation Review and Quality Control:** Our team encountered unexpected obstacles during the review phase of the annotations. Notably, the quality of the annotations took a hit post-review, introducing a significant issue that needed immediate attention. This problem was marked by the appearance of dots on the annotated masks, which compromised the clarity and accuracy of the annotations.

  The introduction of these dots during the review conducted by the client added an extra layer of complexity to the process. It was not just about ensuring the quality of the annotations, but also about dealing with these unanticipated dots that had appeared on the annotated masks.

Detecting Fracture Types

The initial plan for determining the type of fracture (brittle or ductile) involved using the features predicted by the segmentation model. However, we faced a major obstacle with this approach. Given the slight similarities in the features of brittle and ductile images, it became impossible to accurately classify the fracture type using only the segmentation mask. This issue necessitated a change in our methodology.

Direction Marking Logic Adjustment

The plan to mark the direction in brittle fracture images involved tracing a path from the origin to the hackles. However, the team soon realized that this approach had a significant flaw: in nearly 50% of the images, the origin was not present. This absence posed a considerable challenge to our direction marking strategy, as it was not feasible to trace a path from an origin that was not visible in the image.

Model Training

- **Challenges with AWS Sagemaker Notebook Instance Dependencies:** During the model training phase, our team encountered several hurdles related to dependency issues within the AWS Sagemaker Notebook Instance. This cloud-based solution, while robust and flexible, posed certain challenges in terms of managing its dependencies.

  These dependency issues were not minor roadblocks; they significantly impacted the smooth execution of model training. The complexities associated with managing these dependencies within the AWS Sagemaker Notebook Instance required a considerable amount of time, effort, and expertise from our team.

- **Challenges in Selecting Training Categories:** We faced a considerable challenge when it came to selecting the primary categories for model training. From an extensive list of 27 categories, the team had to narrow down its focus to the top 10-13. This task was not straightforward, as it required a careful evaluation of multiple aspects.

Strict Project Timeline

The timeline for this project was indeed stringent. The Daffodil team had a total of

The Solution

eight weeks to complete the project, out of which one week was dedicated solely to User Acceptance Testing (UAT). This left us with merely seven weeks for the actual development phase, a timeline that necessitated intense focus and dedication from the entire team.

THE ROAD TO SUCCESS

Here's how we built a full-fledged solution for the client:

Developing a user-friendly web portal:

In order to further ease the analysis process, a feature-rich, user-friendly web platform / interface / application was designed for our client that allowed analysts to easily upload, analyze and categorize fractures based on microscopic images. Primarily, ReactJs was used to build the front end, while Python was leveraged for the back-end.

Daffodil Software utilized JIRA for project tracking and SCRUM as the software development framework. All project-related documents were uploaded to Confluence, and Bitbucket was employed for code collaboration and version control.

The portal/panel allowed its users to view past reports, review drafts, and accept/ reject pending defects; ensuring a seamless, quick and accurate result on part failure.

In addition, Daffodil also introduced a unique, automated report generation feature in the panel. This was designed to generate defect reports in PPT or PDF formats, providing a unique & versatile solution for documentation and presentation of analysis results.

This was designed to swiftly and accurately identify defect types and propagation direction, thereby eliminating the inconsistencies associated with manual assessment.

Tackling Limited Data Challenge:

To counter this issue, we found ourselves needing to perform extensive hyperparameter tuning. This process involved adjusting various parameters within the model to try and improve its performance despite the limited data. While this was a challenging task, it was necessary to ensure that we could still achieve our goals with the data at our disposal.

The client's data was meticulously sorted into various categories, each corresponding to distinct parameters. These categories were based on the type of material (plastic or rubber), the degree of magnification (10x, 20x, 30x, 50x, 100x, 150x, to 200x), the specific grade of the material (PP, PC, PC PBT, PA ABS, etc.), and the type of fracture (brittle or ductile).

Utilizing Amazon Web Services' Ground Truth

One of the primary reasons behind this decision was the strict compliance requirements stipulated by the client. The client had expressed a clear preference for their data not to be transported out of their existing environment. Given these constraints, AWS Ground Truth emerged as the most suitable option. It allowed us to perform the necessary data annotation tasks while adhering to the client's data security and privacy guidelines. This decision, however, was a tough one, as the team had to quickly familiarize themselves with this new tool and its functionalities.

Improvements in Ductile Image Annotation:

Our initial approach to annotating ductile images involved marking features such as cups and cones, flakes, and stress whitening. However, we encountered several challenges with this methodology. From a machine learning perspective, these features were not distinct enough to facilitate accurate annotations. Predicting directions based on these features proved to be a difficult task.

After facing these hurdles, we engaged in discussions with domain experts to find a solution. These conversations led us to redefine our approach to the annotation of ductile images. We decided to introduce new features that would be more distinctive and thus more conducive to machine learning applications.

The new features we introduced were the high ductility zone, low ductility zone, and stress whitening. These features were more easily distinguishable, making the annotation process more straightforward and accurate. Furthermore, these new features made it feasible to predict directions, an aspect that was previously challenging with the original features.

This enhancement in the ductile image annotation process was a significant step forward. It not only improved the accuracy of the annotations but also made it possible to predict directions effectively, thereby increasing the overall efficiency and effectiveness of our machine learning application.

Maintaining Annotation Review and Quality Control:

The introduction of these dots had a direct impact on our project timeline. The need for a re-review meant additional time and resources had to be allocated to ensure the quality and accuracy of the annotations. The process of rectifying these issues and re-reviewing the annotations was time-consuming. Despite these challenges, our team remained committed to maintaining the highest quality standards for the annotations.

AWS Sagemaker Notebook Instance Dependencies:

To overcome the dependency issues that were hindering the project’s progress, the Daffodil team decided to create a custom UNet model. This model was designed to be compatible with the AWS Sagemaker Notebook Instance and to function effectively despite the dependency challenges.

This custom UNet model not only resolved the dependency issues but also provided the team with a model that was tailored to our specific needs and requirements. While the process of creating this custom model was demanding, it ultimately enhanced the model training process and contributed significantly to the overall success of the project.

Despite the difficulties encountered, our team's ability to adapt and innovate under challenging circumstances was a key factor in the successful progression of our model training process.

Selection of Training Categories:

Firstly, the team had to consider the data count for each category. Categories with a higher volume of data were more likely to provide a robust training base for the model. Secondly, we took into account the features available in the images associated with each category. Categories that displayed more distinctive and easily identifiable features were deemed more suitable for model training.

Additionally, the team had to factor-in the client's needs and preferences. We had to prioritize those categories that were most commonly encountered by the client in their fracture detection operations. This step was crucial in ensuring that the model would be effectively trained to recognize and handle the types of fractures that the client most frequently dealt with.

This multifaceted selection process was challenging, but it was crucial in ensuring that the model would be optimally trained to deliver accurate and useful results for the client

Accurate Detection of Fracture Types:

To overcome limitations of the segmentation model, we decided to introduce classification models into the workflow to overcome this hurdle. These models were aimed at accurately distinguishing between brittle and ductile fractures, thereby improving the overall precision of our project. However, this solution presented another challenge: the timeline.

Training the classification models within this short time frame was a daunting task. Despite this, our team stepped up to the challenge. They dedicated themselves fully to training the classification models.

Their hard work and dedication paid off, and we were able to successfully train the classification models within the limited timeframe. It also significantly improved our ability to accurately detect and classify fracture types, enhancing the overall effectiveness of our project.

Zeroing on the Right AI Models:

Segmentation & classification models, known for their exceptional feature localization, were considered for this project - as other models, despite their strengths, could not provide the fine-grained localization required for this particular task. Moreover, the project demanded a model capable of handling complex and non-uniform feature detection, a task that could pose challenges for other methods.

- **Classification Models:** These models were utilized to categorize images into either brittle or ductile fracture types. Subsequently, based on this classification, the images were routed to their respective segmentation models. Notably, separate segmentation models were employed for brittle and ductile fracture types.
- **Segmentation Models:** Our application of segmentation models aimed to identify and delineate specific features such as Hackles, Mirror Zone (Origin), and others for brittle fractures. Similarly, these models were used to detect features like stress whitening, low ductility, and high ductility for ductile fractures. The rationale behind choosing segmentation models lies in the necessity to precisely extract and highlight these specific features within the images.

Direction Marking Logic Adjustment:

Instead of attempting to trace a path from the non-existent origin to the hackles, we decided to reverse our approach. We would start from the hackles and back propagate to find the origin.

This revised logic for direction marking proved to be a more effective strategy. It allowed us to determine the direction in the images, even in the absence of a visible origin. This adjustment not only solved the challenge the team was facing but also improved the overall accuracy of our direction marking process.

Thorough Documentation Throughout the Project Execution:

During the entire duration of the project, there was a crucial need for us to maintain meticulous documentation. This was in line with the stringent requirements of the Capability Maturity Model Integration (CMMI) Level 5 standards, which demand a high level of process consistency and optimization.

This comprehensive documentation process was not just about maintaining a record of our actions. It was about ensuring that every step we took was in accordance with the highest standards of process maturity and quality. This adherence to CMMI Level 5 standards played a vital role in enhancing the overall quality of our project and ensuring that all processes were optimized and consistent.

Over the course of two months, we produced an impressive total of 43 CMMI Level 5 documents. This vast body of documentation stands as a testament to our commitment to process excellence and quality. Despite the considerable effort and time required to maintain such detailed documentation, we remained dedicated to this task, recognizing its importance in upholding the standards of CMMI Level 5 and ensuring the success of our project.

Managing Strict Project Timeline:

Our team went above and beyond their regular working schedules to successfully ensure that all tasks were completed within the tight deadline of 2 months. This level of dedication was maintained throughout the project duration, demonstrating the team's unwavering commitment to the project's success.

The Impact

Despite the hurdles, the project marked several significant accomplishments.

- Originally intended as a Proof of Concept (POC), the project exceeded expectations by evolving into a full-scale live I-Sense project. It was eventually deployed in the client's production environment, marking a major step forward.
- The team operated effectively in a highly regulated environment and successfully finished the documentation required for CMMI Level 5 compliance.
- Exemplifying exceptional dedication and proficiency, Team Daffodil successfully navigated the challenges and completed the project within the rigorous two-month deadline, thereby reinforcing our ability for timely and efficient execution.
- Along with the team's regular responsibilities, we managed to complete extensive documentation for the project within the set timeline. This comprehensive record played a crucial role in tracking the project's progress and decisions made.
- The impact of the AI solution was profound. The solution not only dramatically decreased the time invested in analyzing, detecting, and reporting failures, but also enhanced the accuracy and consistency of these analyses.
- Daffodil was able to save a total of 60 hours per month (approximately 6.7 man-days) for 30 defects [i.e 2 hours per defect]. When extrapolated, the time savings become even more substantial, freeing up valuable resources that can be allocated to other important tasks.
- Plus, the speed of the AI model's predictions was improved manifolds. It was able to make accurate failure mode predictions in just 2.5 to 3 seconds.

- Lastly, but most importantly, the tool achieved an accuracy rate of 88% in failure mode detection. This high level of accuracy ensured that the results were reliable and could be used to make informed decisions about defect management and prevention.

The success of this POC led to new opportunities. As a result of Daffodil's intense hardwork, the team was able to secure additional projects to work on other Artificial Intelligence/Machine Learning (AI/ML) use cases for the client. This achievement highlighted the success of our project and opened the door for further collaborations in the AI/ML field.

2 hours	Time Saved Per Defect Analysis
2.5- 3 sec	Model Prediction Time
88% accuracy	Failure Mode Detection

Services Used

Computer Vision

AI Development Services

Software Product Engineering