Pega Process AI Essentials

STUDENT GUIDE

PEGA® Academy

This document is the property of:
Pegasystems Inc.
1 Rogers Street
Cambridge, MA 02142
Phone: (617) 374-9600
Fax: (617) 374-9620
www.pega.com

Mission: Pega Process AI Essentials

Product: Pega Platform™ '24.1

URL: https://academy.pega.com/mission/pega-process-ai-essentials/v5

Date: 04 April 2024

Contents

Pega Process AI overview

Description

Gain a greater understanding of the key features, capabilities, and benefits of Prediction Studio in Pega Process AI™ context. Prediction Studio is the dedicated workspace for data scientists to control the life cycles of predictions and the predictive models that drive them. Configure the predictions that are deployed in Pega Process AI to increase efficiency and effectiveness in case management.

Learning objectives

- Describe the use of Pega Process AI in case management
- Explain the types of predictions that are available in Prediction Studio

Pega Process AI overview

Introduction

In recent years, artificial intelligence has moved out of the labs and helped enterprises generate proven business value. At the same time, operationalizing AI can be a bottleneck. Pega Process AI™ tackles this problem by using AI to self-optimize processes and applying your own AI in Pega case management.

Transcript

This video provides an overview of the Pega Process AI capabilities in intelligent automation.

Process management aims to optimize business processes by increasing efficiency, consistency, and transparency, which decreases costs and improves quality.

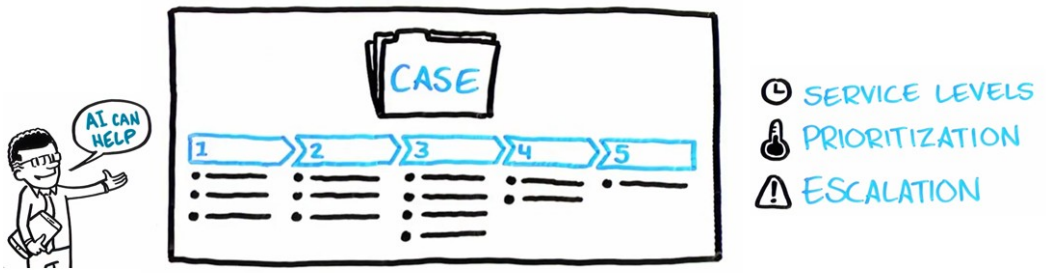

For example, consider an online order process. The customer submits an order, and the company processes and then delivers the order.

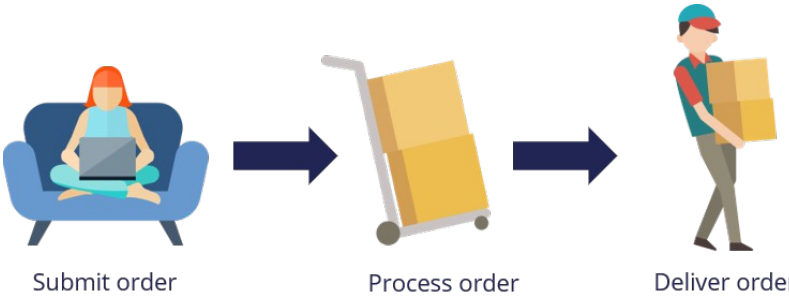

An Pega Platform application that models the online order process follows the same sequence as a series of stages. A **case type** is the abstract model of that process.

Case types model repeatable business transactions that might refer to a customer, or another entity, such as a machine in a maintenance case type.

The **case life cycle** for a case type helps to visualize the work to complete as part of a business transaction.

Each stage in the life cycle contains the steps required to complete it and move to the next stage.

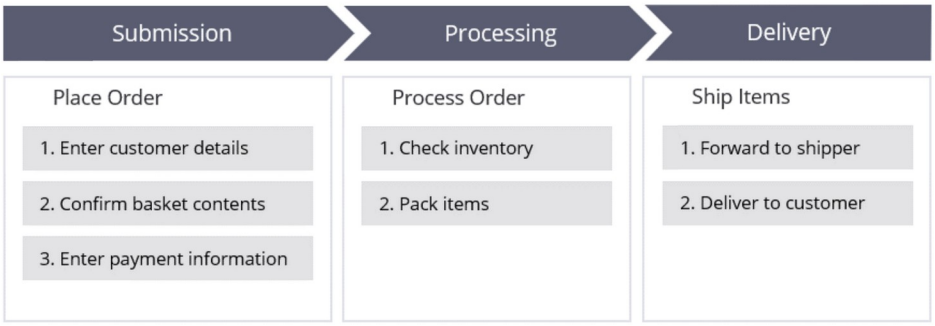

A **case** is a specific transaction instance of the case type.

Each time a user submits an online order, Pega Platform creates an order case and assigns the case a unique identifier.

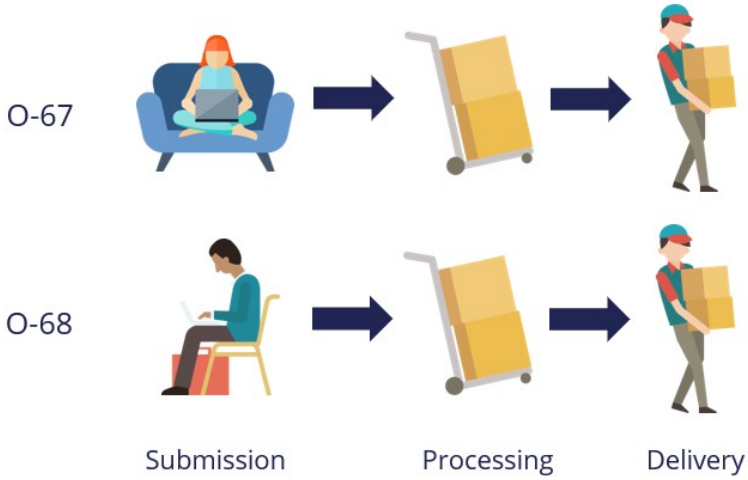

A case type can use declarative rules to manage the workflow, for example, to confirm that the order contains a valid shipping address or the order amount threshold to qualify for free shipping.

Pega Process AI can improve the quality of the decisions in the workflow by weighting in predictions, driven by predictive models.

The first approach is to operationalize existing predictive models that have proven their efficiency, to support the decisions that benefit from predictions, such as credit risk in a sales case or fraud risk in a claims case.

For example, the decision not to process an order can be based on a high credit risk score, and then the application can route the dubious claim for closer inspection.

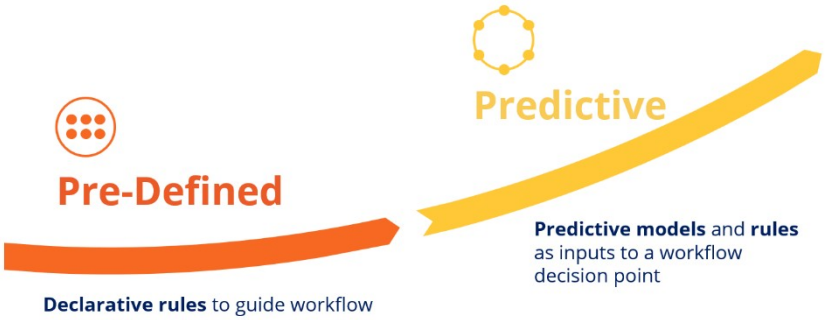

The inputs for such a predictive model can be attributes of the case itself, such as the claimed amount in a claims case type, but they can also include data such as the number of claims submitted recently by the same customer.

You can build predictive models in Prediction Studio, import the models in the PMML and H2O formats, or run externally on the Amazon SageMaker and Google ML platforms to drive a prediction.

To optimize case outcomes, use adaptive models that can predict outcomes, such as case resolution, or intelligently prioritize and route cases to optimize business value and customer experience.

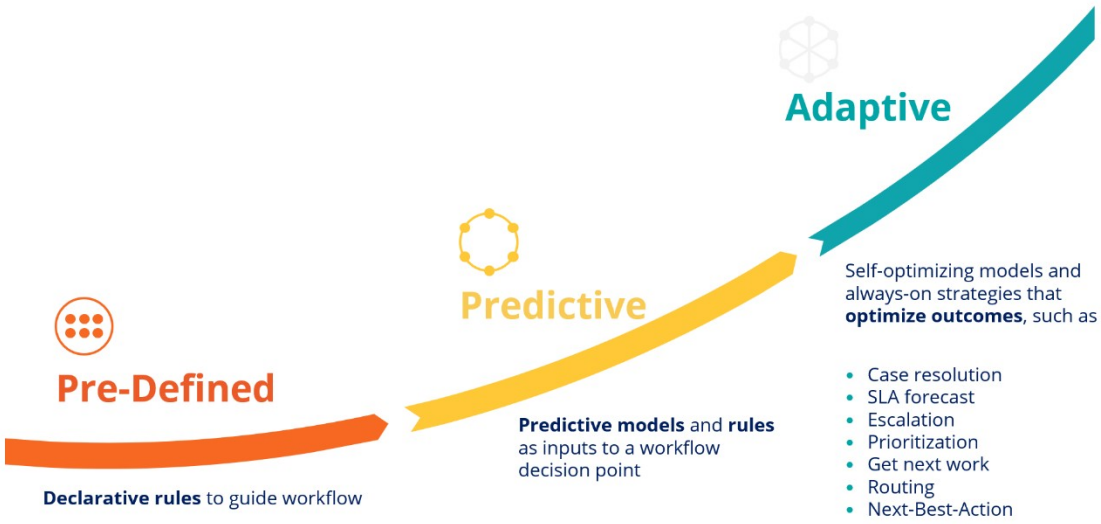

Adaptive models self-optimize by learning from the previous case outcomes that they capture.

The objective of Pega Process AI is to make sense of the incoming data and then decide on the best action to take in a specific stage of the case.

You can enhance the incoming data analysis by event processing to detect patterns of interest in real-time data streams and by natural language processing of incoming text.

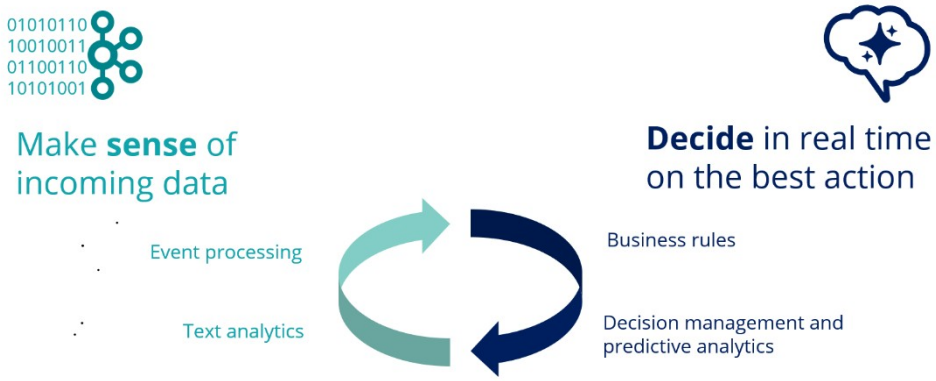

The decision is based on the business rules and supported by predictive analytics. This process is repeated every time that the case requests a decision.

As the number of processed cases increases and model evidence accumulates, the predictive power of the models increases over time.

To summarize, Pega Process AI uses artificial intelligence in case management to produce better business outcomes.

You can use real-time, adaptive case outcome predictions and your own AI models in custom predictions.

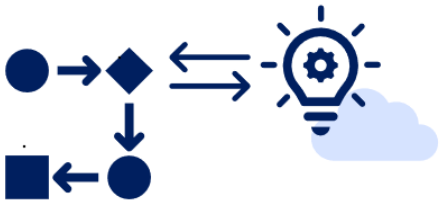

Custom AI predictions

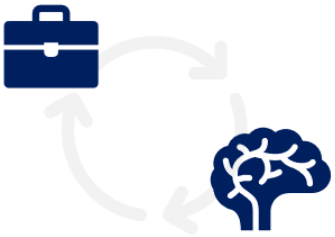

Real-time, adaptive **case outcome predictions**

Process AI predictions

Introduction

With the decision management capability of Pega Platform™, you can enhance applications to help optimize business processes, predict customer behavior, analyze natural language, and make informed decisions to better meet the needs of customers and achieve positive business outcomes.

Transcript

This video introduces you to Pega AI, a feature of the decision management capability of Pega Platform.

Other decisioning features of Pega Platform include:

- Decision strategies that feature a business- and user-friendly canvas with which you can create decision logic that uses behavioral and operational data to improve intelligent processes.
- Event strategies to detect patterns in data streams and react to them.
- Data flows as scalable and resilient data pipelines to ingest, process, and move data from one or more sources to one or more destinations.

Decision management uses Pega AI to make predictions about the possibility of fraud, successful case completion, and other subjects to make decisions more relevant.

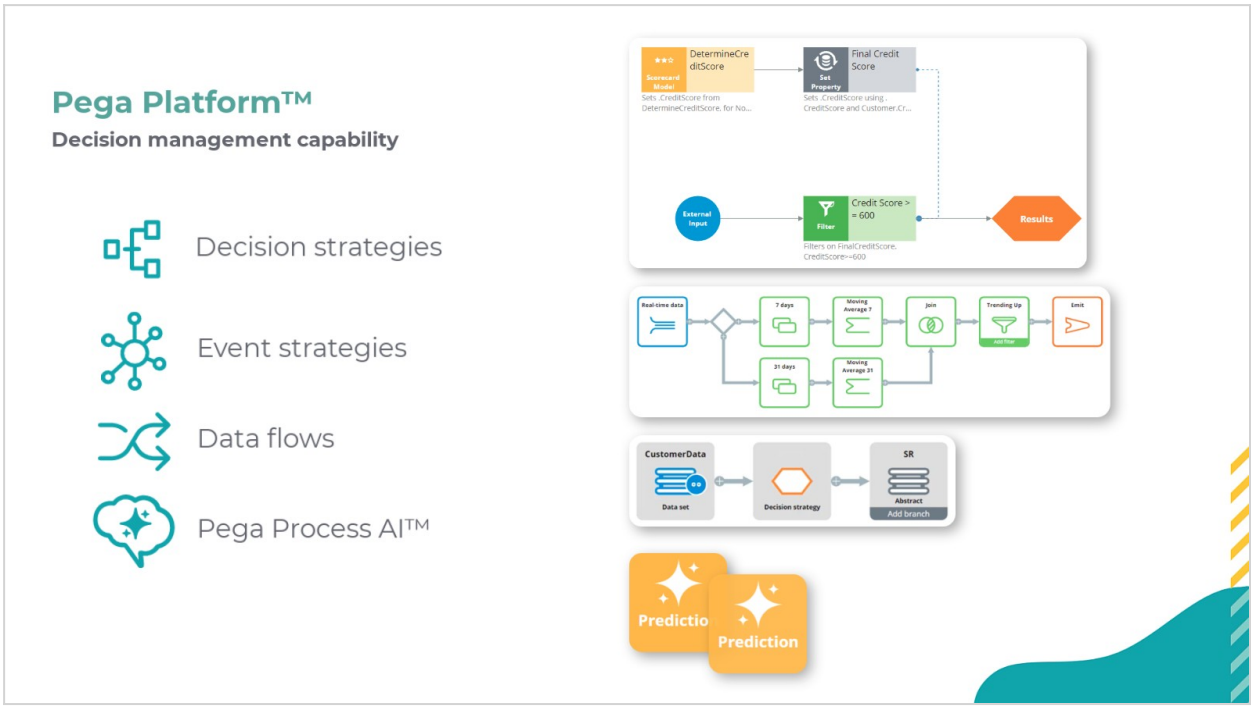

Decision management is a Pega Platform capability. You can apply decision management to any application that is built on Pega Platform.

Various and versatile predictions are available, but one or more predictive models drive them all.

For example, a data scientist can create a predictive model in Pega Platform or an external environment that can export the model as a PMML or H2O file back to Pega Platform. Another option is to connect to a machine learning service such as Google ML or AWS SageMaker.

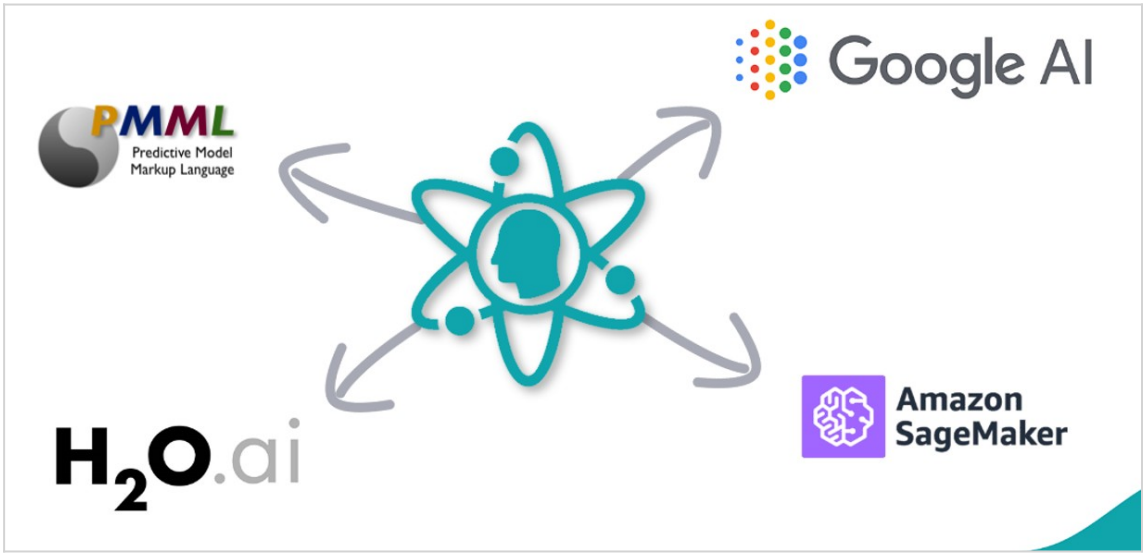

If an insurance company wants to use Pega Process AI™ to route incoming claims that might be fraudulent to an expert, based on the outcome of a predictive model, the data scientist creates a fraud model to drive a new case management prediction in Prediction Studio.

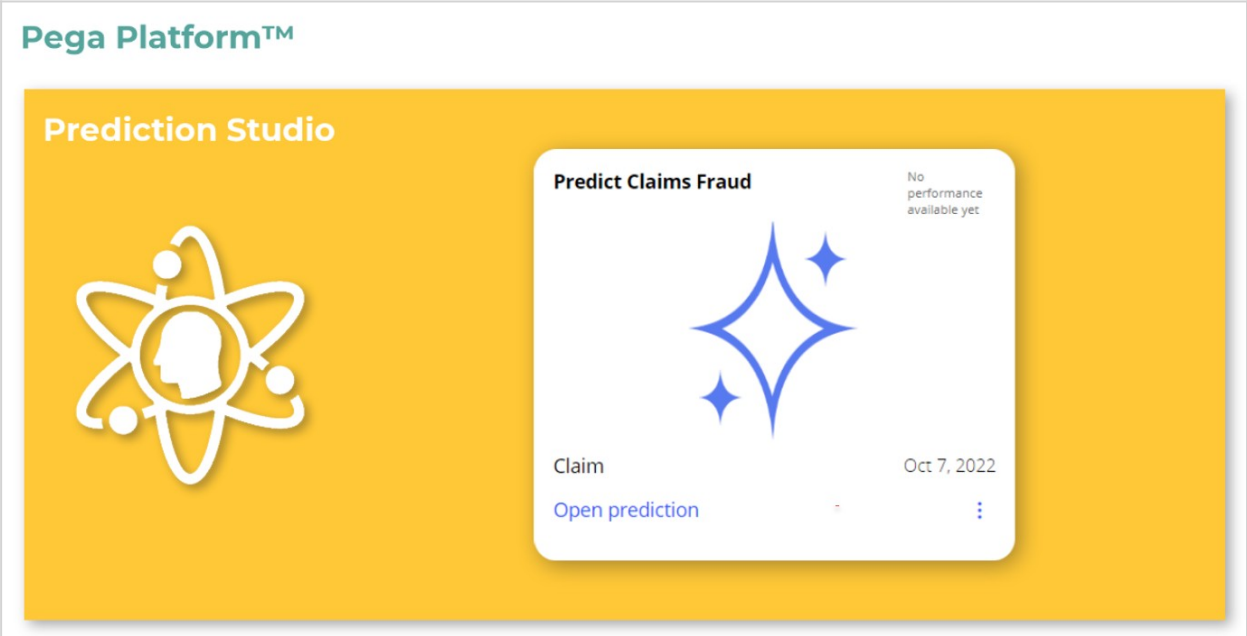

Prediction Studio is the dedicated workspace where you manage the life cycle of predictive models and the predictions that they drive.

The workspace provides data scientists with everything they need to author, deploy, govern, monitor, and change predictions. Prediction Studio has five work areas: **Predictions**, **Models**, **Data**, **Reports**, and **Settings**.

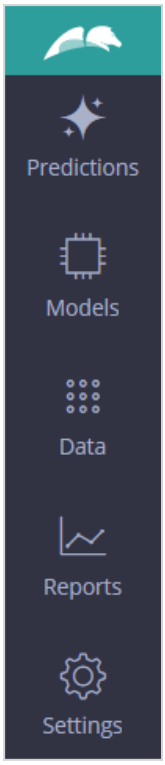

On the Predictions landing page, you create and manage predictions. There are three types of predictions, but Process AI focuses on the **Case management** prediction.

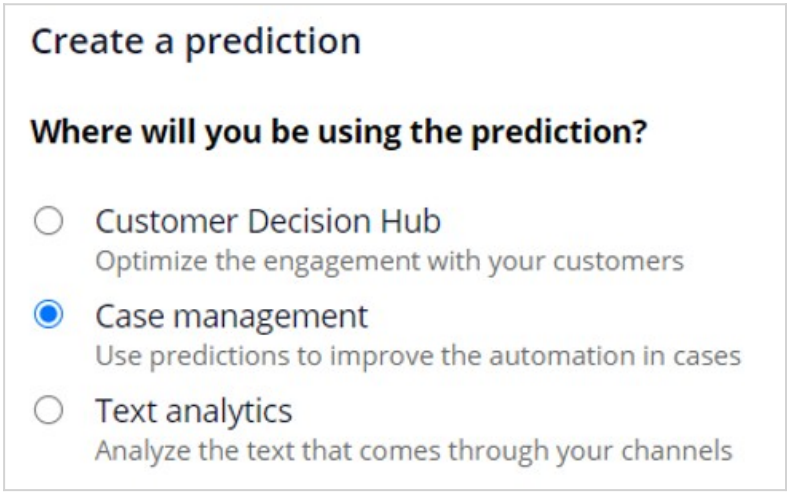

Case management predictions are used in case types to support decisions in business processes. For example, predictive models can help to predict whether an insurance claim is fraudulent or distinguish regular from complex claims.

This dependence routes cases more accurately and strengthens the separation of concerns.

The decision step in a case type uses **case management** predictions.

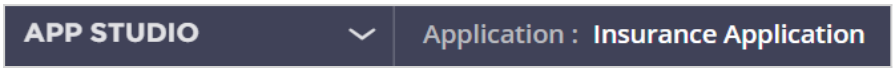

Consider the following case type, which handles incoming car insurance claims:

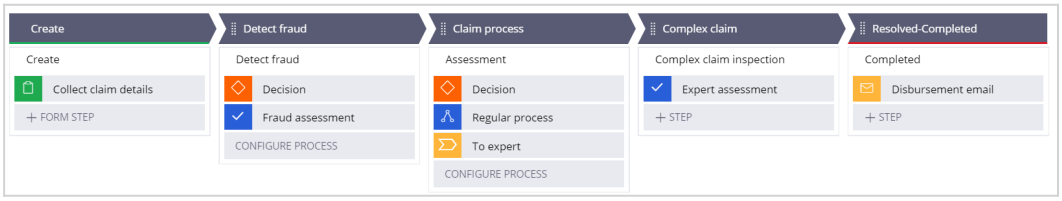

An application developer can use the outcome of the prediction in the condition of a decision step instead of a business rule. Based on the condition, the system routes a case to a fraud expert when the prediction flags the claim as abnormal.

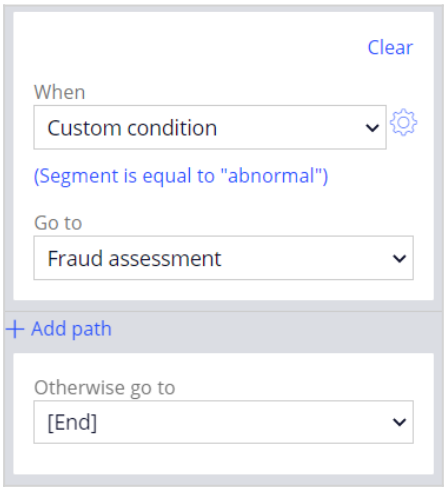

Pega Process AI can also help to distinguish regular from complex claims. It helps speed up the process by identifying such cases early and routing them to the right person.

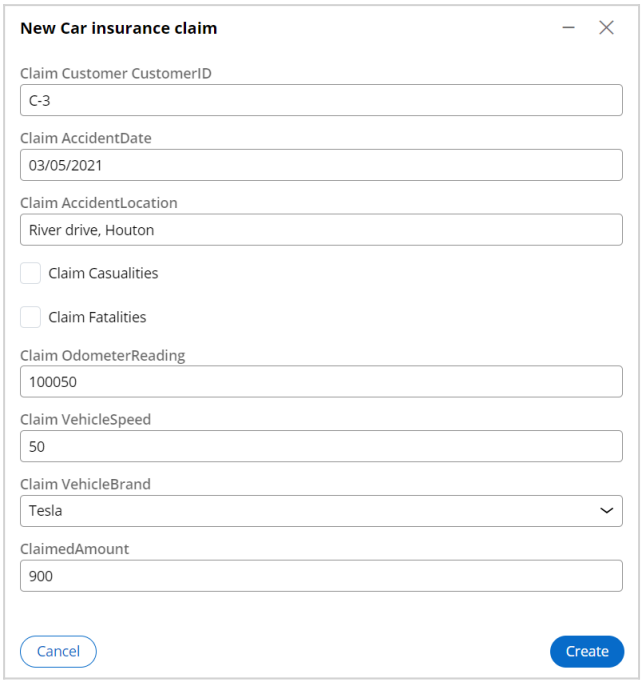

In the following case, the data scientists create a prediction that aims to identify cases that are likely to miss their deadlines. Then, an application developer uses the prediction outcome when configuring the case type so that the system can automatically route a complex case to a senior employee for evaluation.

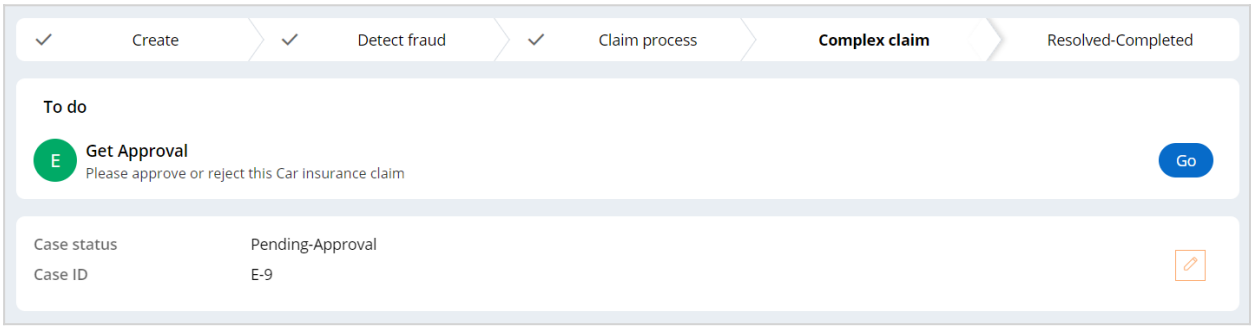

Some additional helpful information widgets are available in the case view. Notice the prediction widget in the claims case. The widget conveniently displays the output of the trained prediction models, in this case, fraud or missing service-level agreement (SLA) probability. By clicking **Learn more**, you can see the details of the model and whether it received training with adequate data (in this example, the probability of case completion).

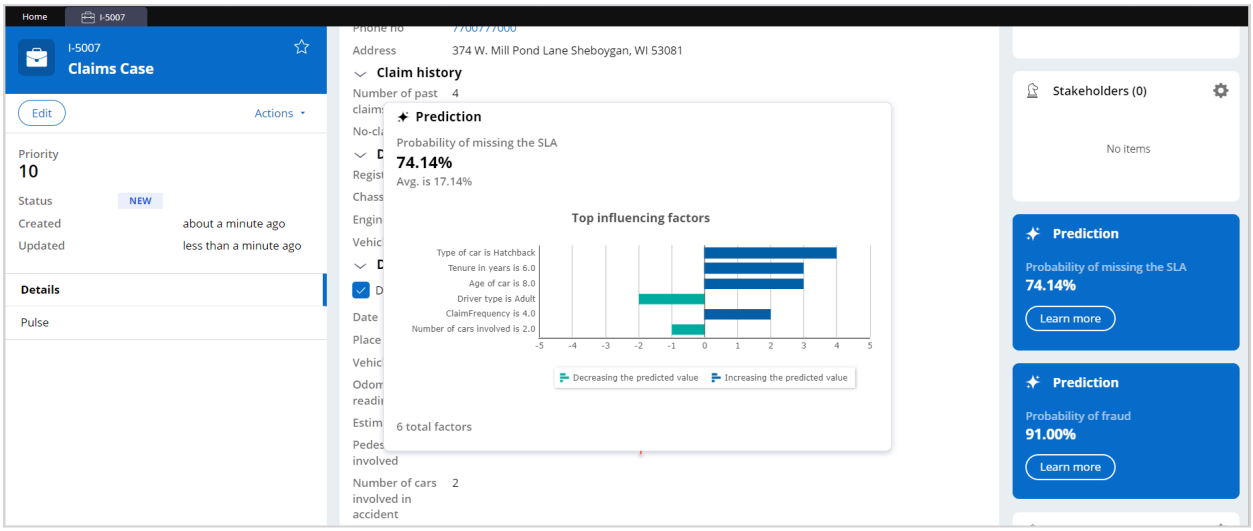

The widget also displays which predictors contribute positively (increasing the predicted value), and which contribute negatively, (decreasing the predicted value).

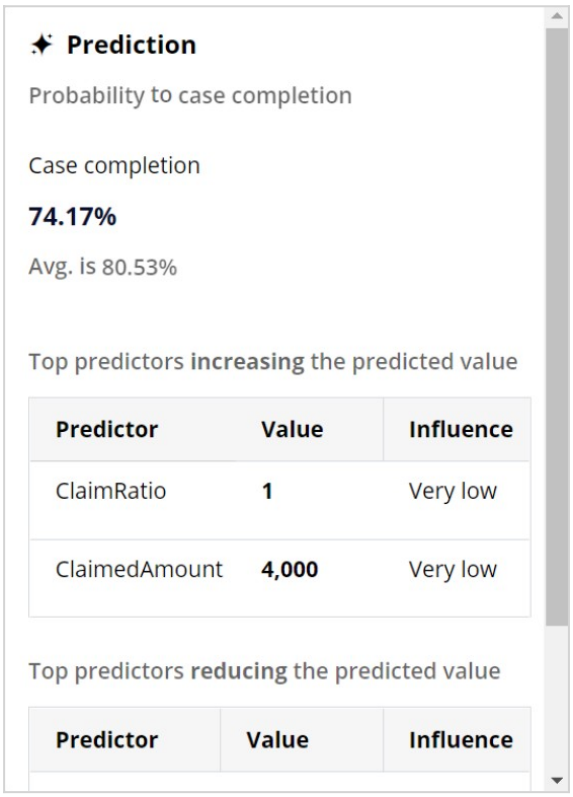

You have reached the end of this video. You have learned:

- How Pega AI allows you to improve business processes by using predictions.
- How predictive models drive predictions.
- How to create and manage predictions in Prediction Studio.

- How to use predictions in a case type to improve business processes.

Applying NLP for Case classification

Description

Pega Process AI utilizes natural language processing to automatically categorize and route claims based on customer descriptions, eliminating delays and errors associated with manual processing. Learn how to enable AI Case classification, define topics, train models, and validate the accuracy of automatic categorization to streamline claims processing and enhance operational efficiency.

Learning objectives

- Enable AI Case classification in Pega Process AI by creating a new accident category field and defining topics.
- Train and test models in Prediction Studio
- Validate the automatic AI accident category prediction and Case routing by creating a new Claims Case for a test Persona.
- Verify the effectiveness of intelligent Case routing, which automatically routes Cases to the correct Work Queue based on the predicted accident category.

Applying NLP for Case classification

In the complex and dynamic world of insurance, processing claims is an important task. It requires a high degree of precision, efficiency, and speed. However, the traditional approach, which involves manual categorization and routing of Cases, often leads to delays and errors. These issues can negatively impact customer satisfaction and operational efficiency.

Pega Process AI, equipped with advanced features, addresses these challenges. By using the power of natural language processing (NLP), Pega Process AI can automatically identify the accident category based on the Case description provided by the customer in their insurance claim. As shown in the following demo, Pega Process AI not only eliminates the need for manual intervention but also ensures swift and accurate routing of the Case to the correct Work Queue.

Transcript

This demo shows you how to enable Case classification with AI and test it in a real-world scenario.

Consider a scenario that involves an insurance company, U+ Insurance. This company is currently handling a high volume of car insurance claims. The process they follow is traditional and manual; experts review all Cases and manually categorize and route them based on the description provided by customers. This process often results in delays and errors due to the sheer volume of Cases and the human element involved.

Recognizing the need for a more efficient and accurate process, the company decides to utilize natural language processing in Pega Process AI.

The goal is to automate the categorization and Case routing, which reduces the time it takes to process claims and minimizes the possibility of errors as a result.

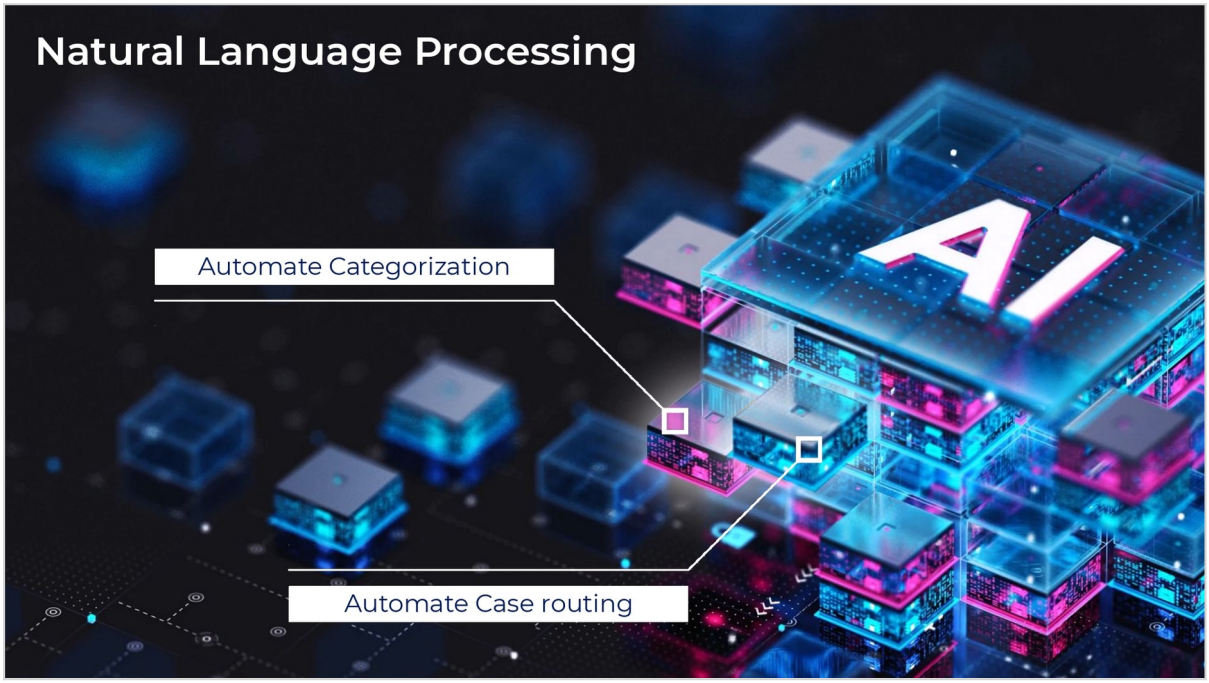

An Application Developer must enable the AI-powered Case classification prediction. If the desired category field is nonexistent in the application, the developer begins by creating it in the Case and defining the topics. The system uses the choices set by the Application

Developer as topics for the new Accident Category Prediction. This is a crucial step as it sets the foundation for the AI to understand and categorize the Cases. Now, the Application Developer must enable AI and specify the Case description provided by the customer as the input text for AI analysis.

After the Application Developer enables AI, the system automatically creates a new Prediction in Prediction Studio. When insurance claims arrive, the Prediction uses AI models to recognize the topic of each claim and categorize it based on the Case description provided by the customer. The topics correspond with the choices set in the new Accident Category field.

At first, the models are not trained. Data Scientist trains and builds the models in Prediction Studio. Data Scientist uses a provided training data file, which contains examples of Case descriptions and their corresponding topics. The topics match accident categories that the

prediction needs to detect.

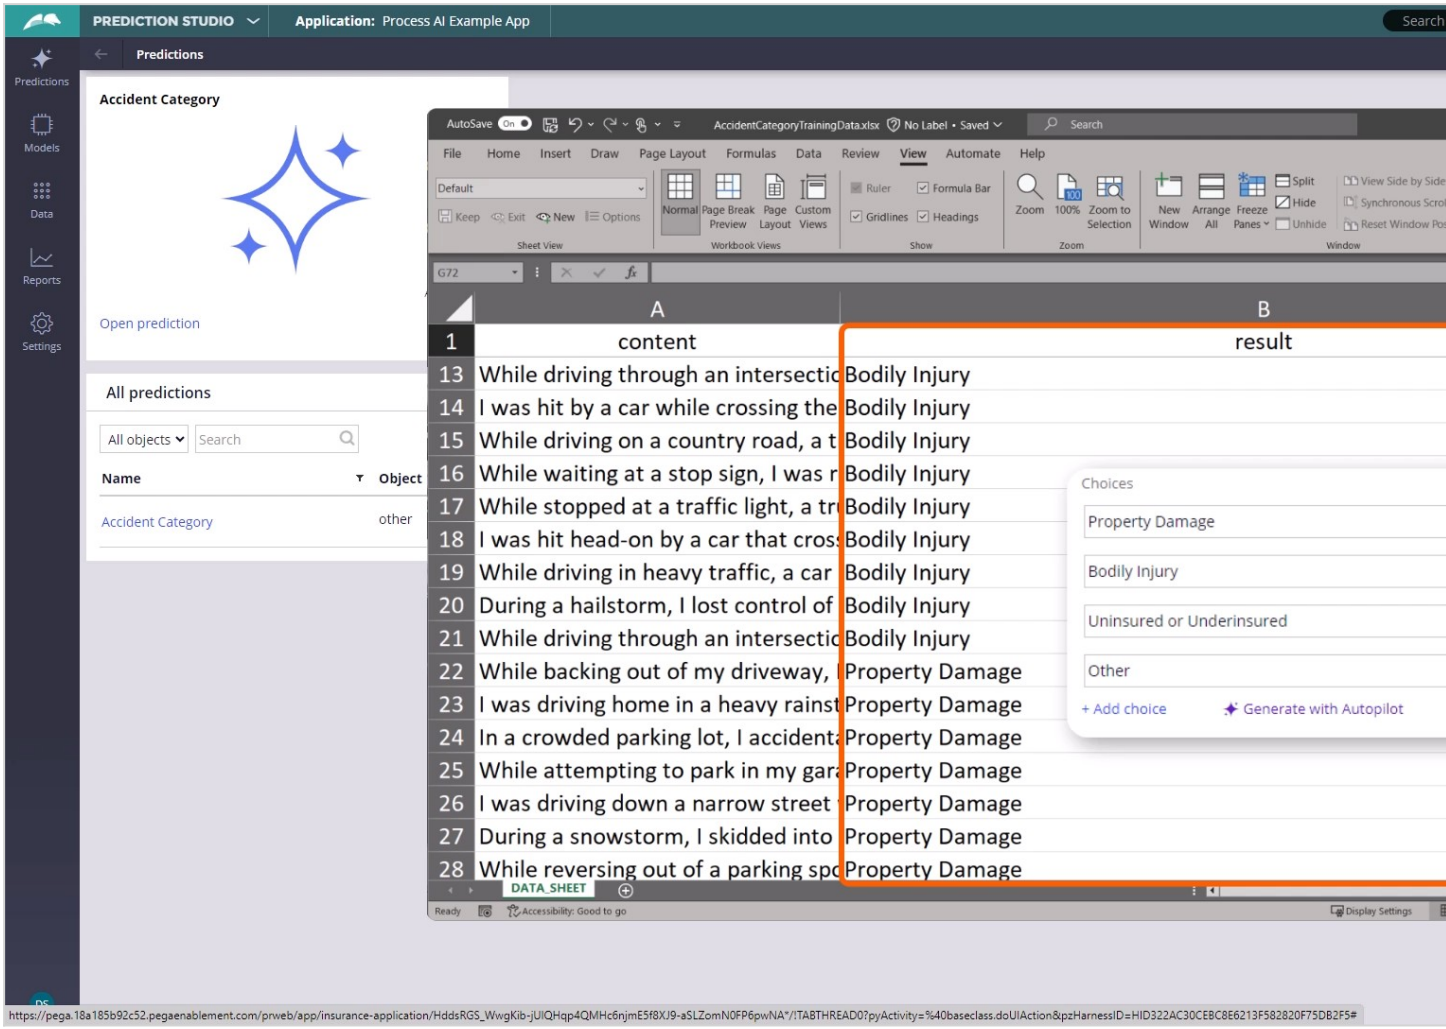

After the models train on the data, the text Prediction undergoes testing in Prediction Studio. The Data Scientist confirms that the AI correctly detects the topic of each message by providing the AI with sample Case descriptions and checking if the AI correctly identifies

the accident category.

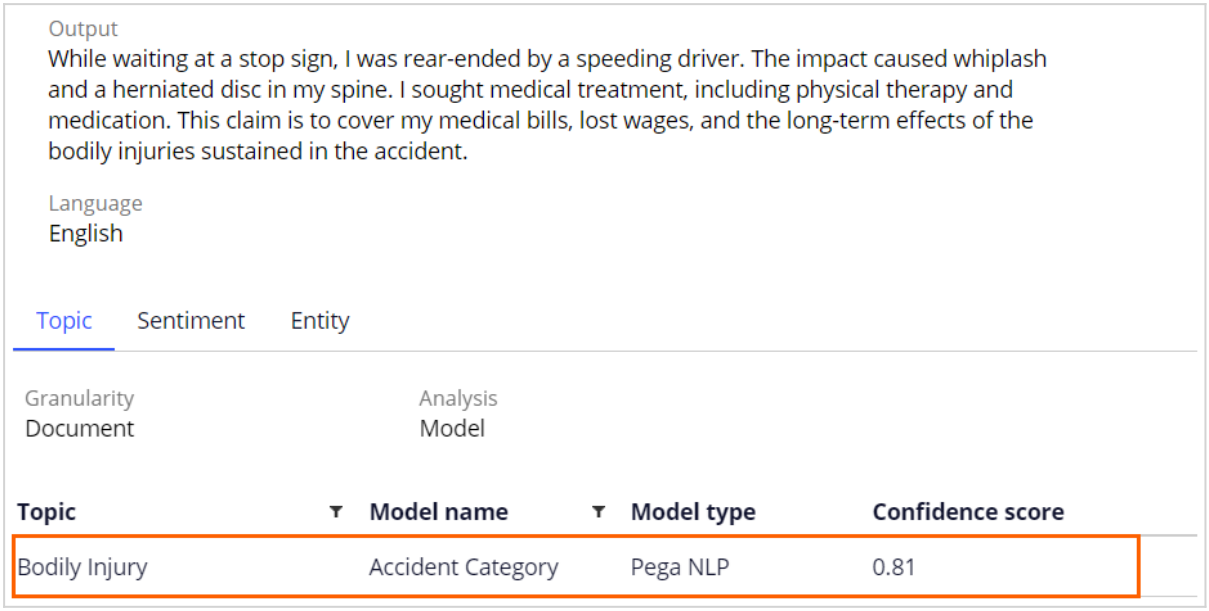

After the models complete training and testing, the Application Developer steps back in to validate the functionality of the automatic AI accident category prediction and Case routing in the Process AI Example Application. The developer creates a new Claims Case for each Persona and validates the detected category. This step helps ensure that the AI is accurately categorizing and routing the Cases in a real-world scenario.

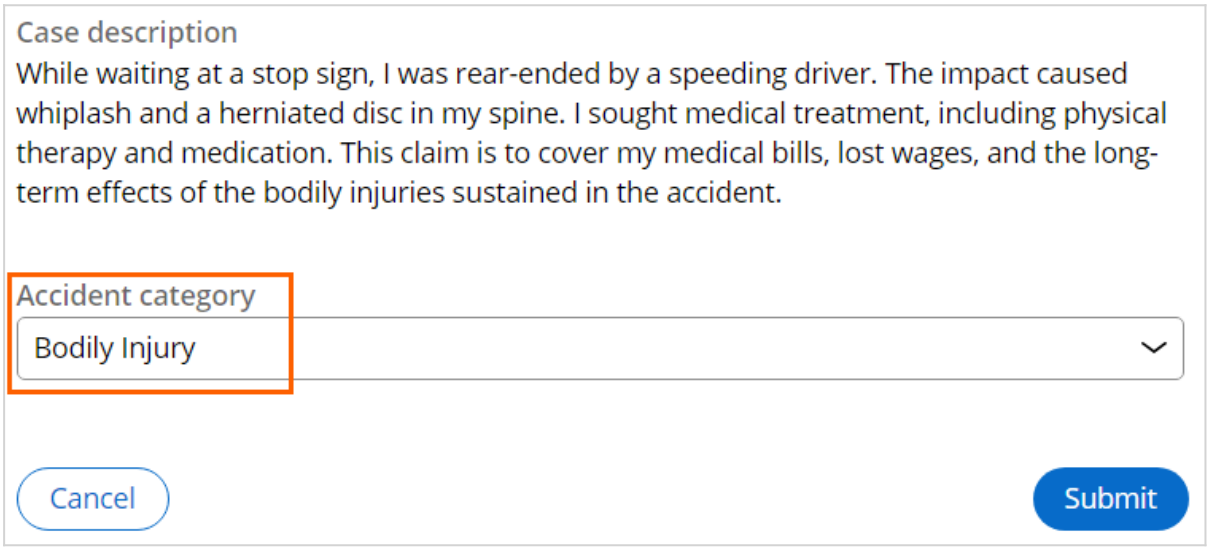

Finally, the Application Developer verifies the pre-configured intelligent Case routing. The system automatically routes the Cases to the correct Work Queue based on the predicted accident category. This action eliminates the need for manual classification and ensures that the appropriate teams handle Cases, which improves the efficiency of the claims

processing workflow.

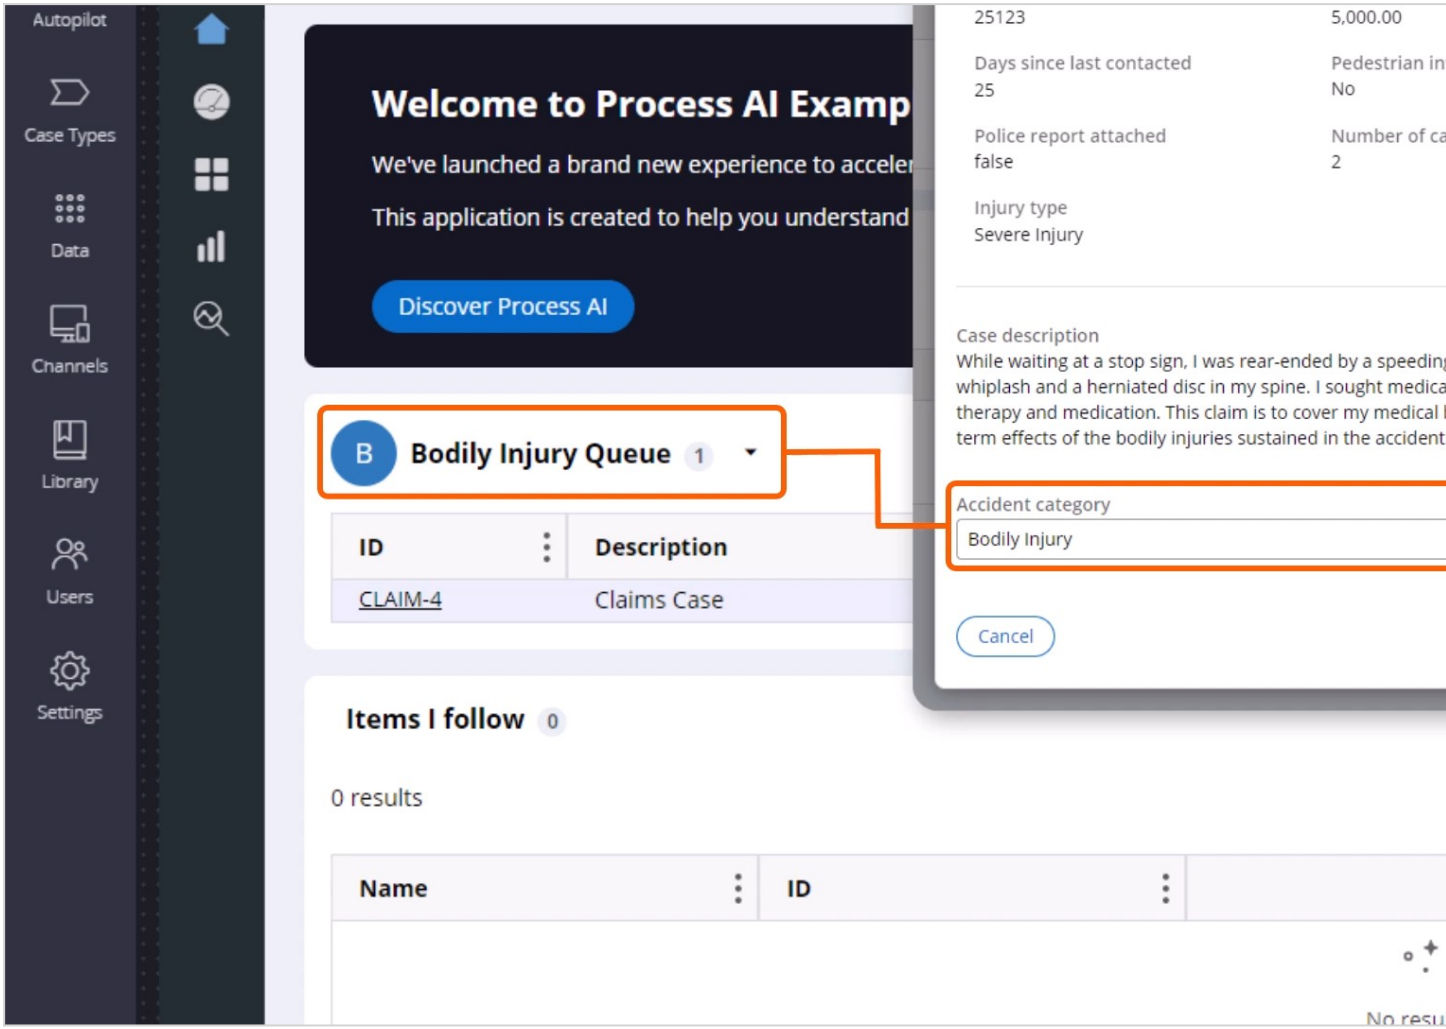

This demo has concluded. What did it show you?

- How to enable AI Case classification in Pega Process AI by creating a new accident category field and defining topics.
- How to train and test models in Prediction Studio.
- How to validate the automatic AI accident category prediction and Case routing by creating a new Claims Case for a test Persona.
- How to verify the effectiveness of intelligent Case routing, which automatically routes Cases to the correct Work Queue based on the predicted accident category.

Predicting fraud

Description

Occasionally, an insurance claim might be erroneous or even fraudulent. To detect fraud and optimize the way in which the application routes work and meets business goals, learn how to use your own predictive models in case management.

Learning objectives

- Create a prediction to detect fraud
- Use the new prediction in a case type

Predicting fraud

Introduction

Pega Process AI™ lets you bring your own predictive models to Pega. Use predictions in case types to optimize the way in which your application processes work and to meet your business goals. Learn how to use a predictive fraud model to effectively route suspicious claims for closer inspection.

Transcript

This demo will show you how to use a predictive fraud model in a case type to route suspicious claims to an expert.

U+ Insurance uses Pega Platform™ for case management. The life cycle of the case type that processes incoming car insurance claims contains a fraud detection stage, a regular process stage, and a complex claim process stage.

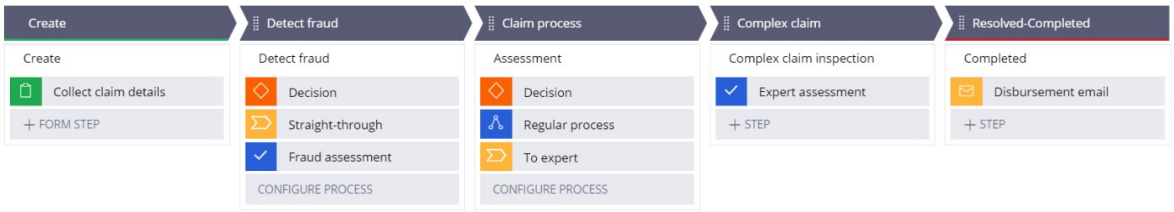

When the case is resolved, the claimant receives an email that communicates the decision.

The decision step in the **Detect fraud** stage routes cases with a low claimed amount for straight-through processing.

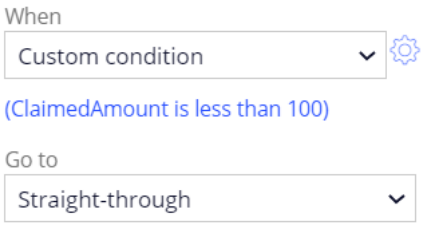

A set percentage of claims with a high claimed amount is routed to an expert for fraud assessment.

Consider this car insurance claim. The claimed amount is 50.

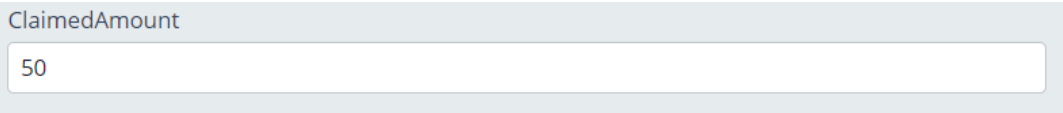

The claim qualifies for straight-through processing as the claimed amount is below the threshold. The case is automatically resolved, and the claimant receives an email that states that the claimed amount will be disbursed.

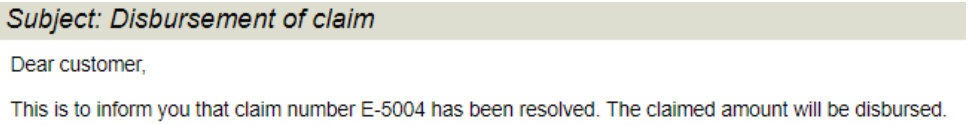

A fraud expert inspects a set percentage of cases with a high claimed amount.

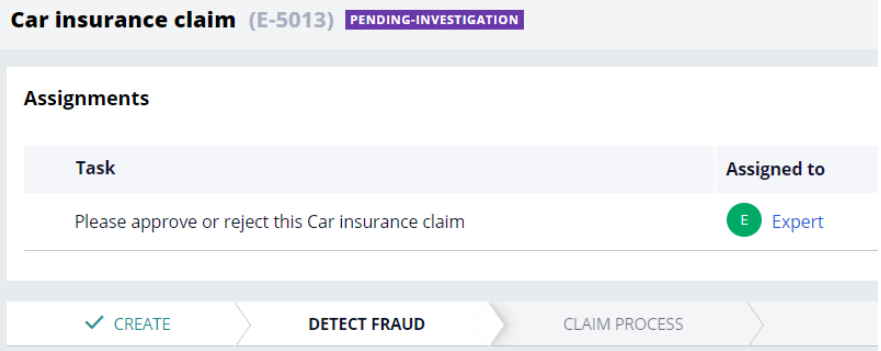

After approval, the system routes the case to the regular claim process.

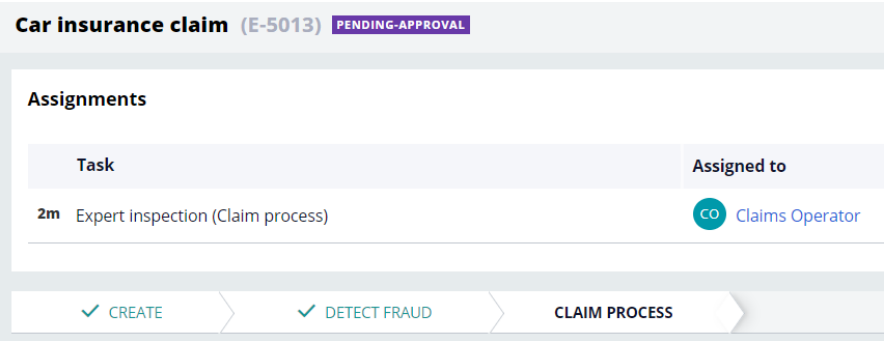

U+ Insurance wants to improve the effectiveness of fraud detection by using a predictive model that calculates the fraud risk of each claim.

The business requirements are that claims only qualify for straight-through processing if the fraud risk score is very low, while all claims with a high fraud risk score are inspected by the fraud expert. The routing of randomly selected cases to the fraud expert must remain in place to create a control group.

The data scientist team of U+ Insurance has developed a fraud model on the H2O.ai platform and has validated the model against historical data that the company captured.

The system qualifies a claim as abnormal if the probability of fraud exceeds the threshold; otherwise, the system classifies the case as normal.

To implement the fraud model, you create a new case management prediction. You can create a custom prediction that can forecast binary or numerical outcomes.

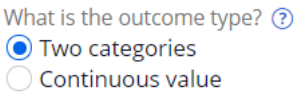

For fraud detection, Process AI provides an out-of-the-box template. The claim is the subject of the prediction.

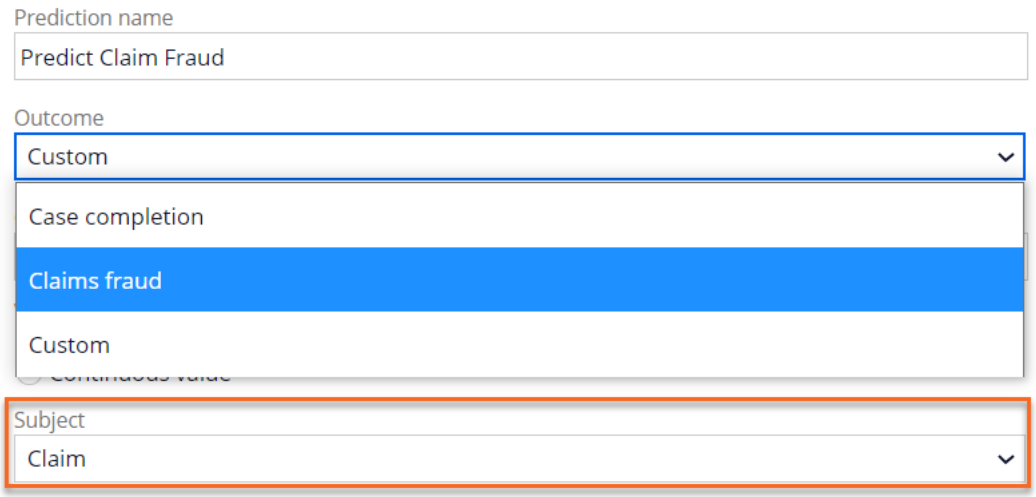

A placeholder scorecard initially drives the prediction.

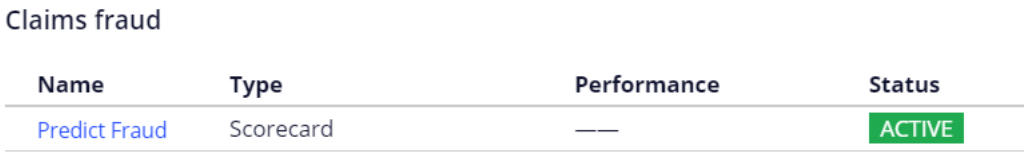

When the predictive fraud model replaces the scorecard, the prediction is ready for implementation in the Car insurance claim case type. You replace the placeholder with a machine learning model, a scorecard, or a field that contains a precalculated score. You can upload a machine learning model as a PMML or H2O file. Alternatively, you can connect to online machine learning services.

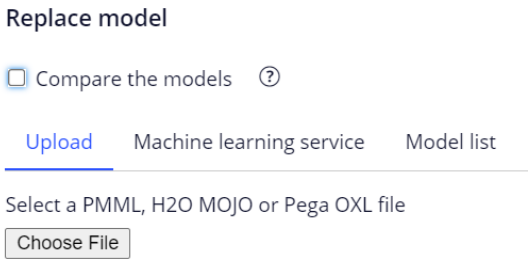

You can select predictive models that are available in the application in the model list. When the model is ready for review, approve the model to replace the scorecard.

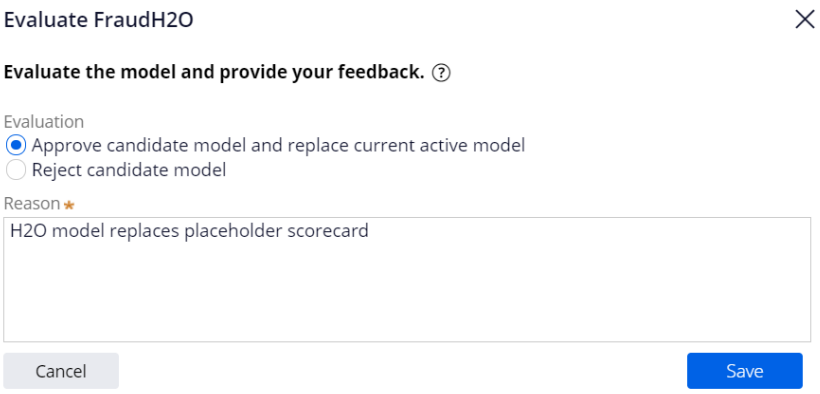

The fraud model now drives the prediction.

When you run the model with these input values, the model qualifies the claim as abnormal.

Field name	Type	Input
EntitledAmount	Double	10000
ClaimedAmount	Double	900
ClaimFrequency	Double	1
SuspiciousClaim	Double	2
ClaimRatio	Double	0.8

The model predicts the claim to be abnormal because the propensity value is above the threshold.

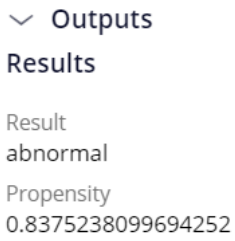

Predictors of the model include the claim data, such as location and claimed amount, but can also cover customer behavior data, such as the number of recent claims.

As an application developer, you can implement the fraud prediction to route claims based on the fraud risk calculated by the model.

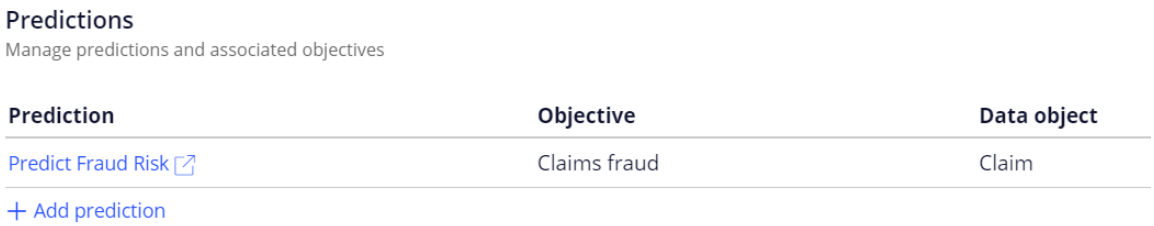

In the **Decision** step in the **Detect fraud** stage of the life cycle, implement the prediction. Add the condition that only claims with a very low predicted fraud risk qualify for straight-through processing.

Replace the condition that routes a claim to a fraud expert based on the claimed amount with a condition that is based on the outcome of the fraud model and change the logical operator to generate the control group.

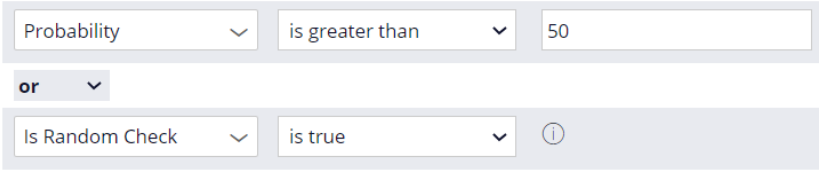

When you run the same claim that previously qualified for straight-through the claim now disqualifies because the condition that fraud risk is very low is not met and the system consequently routes the case for regular processing.

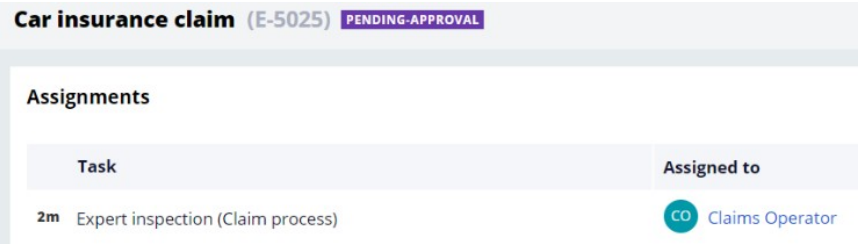

When a claim with the same predictor values as previously tested in Prediction Studio is run, the system routes the case to the fraud expert.

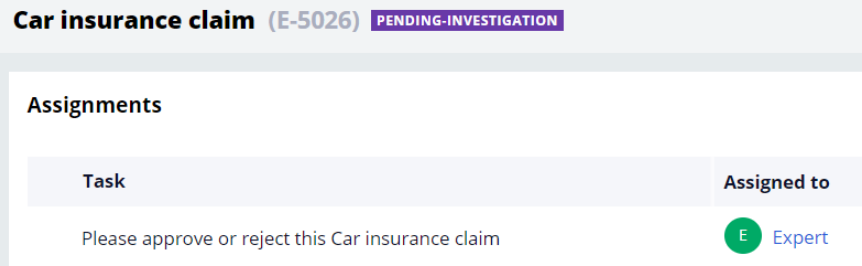

This demo has concluded. What did it show you?

- How to create a case management prediction driven by a predictive model.
- How to use a prediction in a case type.

Predicting missing the Service-Level Agreement

Description

Pega Process AI™ can help to distinguish regular from complex claims. Complex claims often escalate into a lengthy process, which is not only costly, but also leads to poor customer experiences.

Learn how to use Process AI to create an adaptive model to route complex cases to an experienced handler and leave many of the claims for straight-through processing. As the adaptive model learns from the outcome of each case, it becomes more accurate at predicting which claims to escalate, and in that way to self-optimize the process.

Learning objectives

- Create a prediction that predicts case outcomes
- Use the new prediction to route complex cases to an expert

Predicting missing the Service-Level Agreement

Introduction

Pega Process AI™ can help to distinguish regular from complex claims. Complex claims often escalate into a lengthy process, which is costly and leads to a bad customer experience. The distinction lets you detect these claims early and address them at once.

Learn how to create a prediction that aims to identify cases that are likely to miss their deadlines and route them to a senior employee to handle them more efficiently and improve the customer experience.

Transcript

This demo shows you how to use adaptive models to predict missing the Service-Level Agreement (SLA).

U+ Insurance uses Pega Platform™ for case management. An incoming car insurance claim is straight through processed, or routed to a claims operator, who approves or rejects the claim to resolve the case.

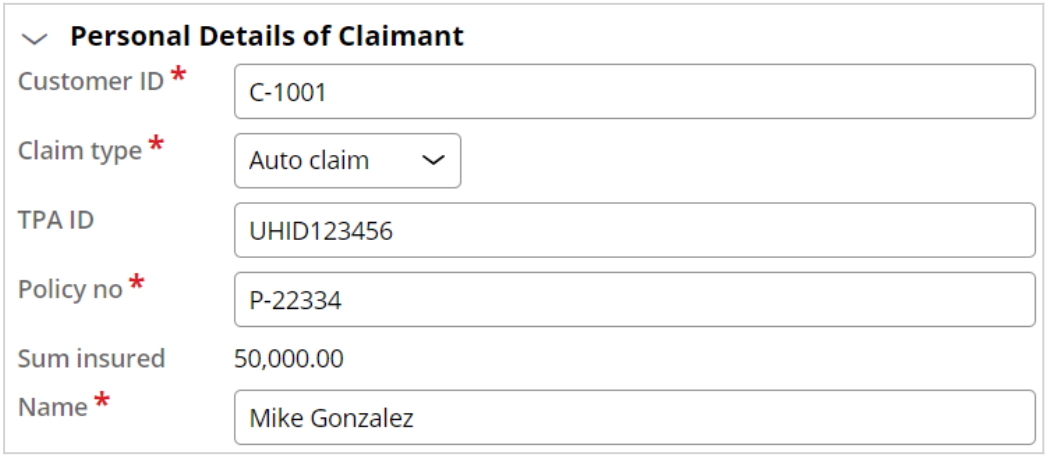

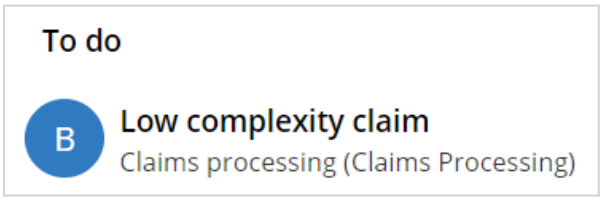

A case is escalated to an expert when the claim is not completed in the allotted time for regular processing.

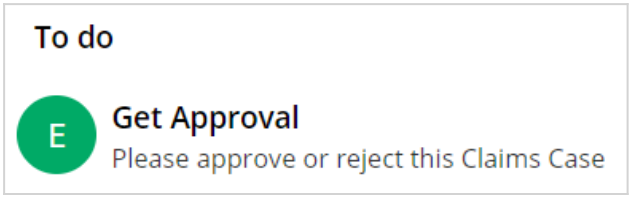

In the current configuration, claims that exceed 45000 are considered highly complex and are always investigated by an expert as a precaution.

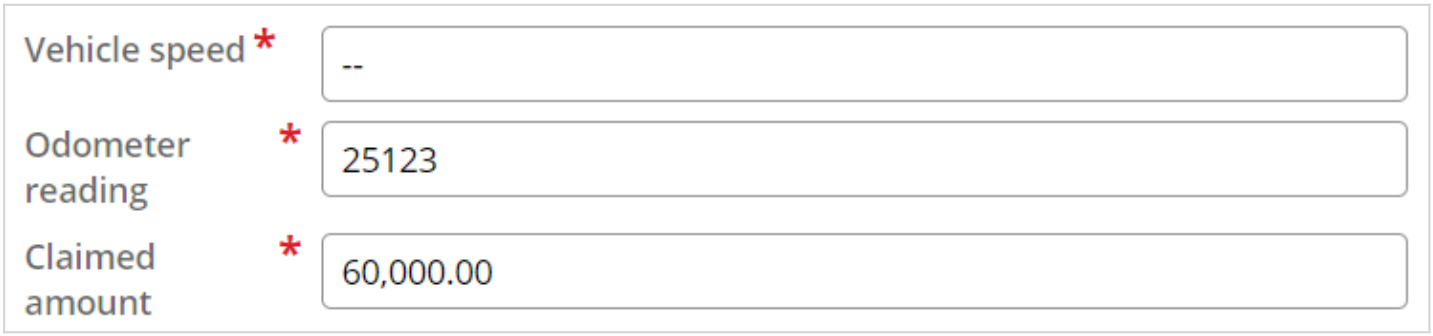

However, decisioning using hard business rules, like in this case based on a simple cutoff value is not efficient, because even cases that exceed 45000 can be often resolved on time in the regular claims process. As a result, the experts consequently spend valuable time on relatively simple claims.

Process AI can help optimize the process by predicting the likelihood that a case is resolved before the deadline in the regular workflow and otherwise, route it to an expert irrespective of the cause of the complexity of the claim. First, it is an Application Developer's task to create a Boolean outcome field. It serves the adaptive model as the outcome field and allows it to distinguish cases that missed the SLA. You add the outcome field in the case type data model settings to make it available in that case type.

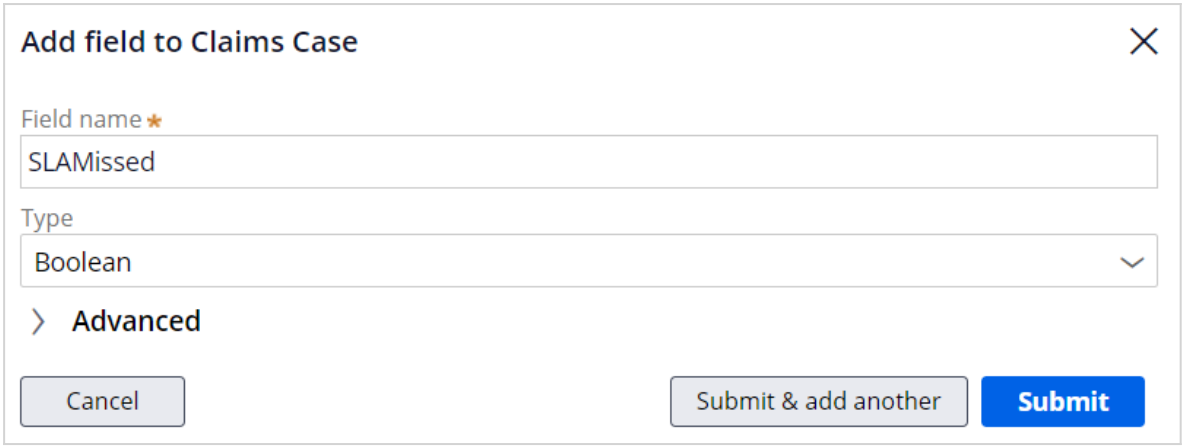

Next, to allow the model to learn from future outcomes, in Goal & deadline settings of the case type, the Application Developer configures a condition that automatically sets the outcome as missed when the deadline expires. Finally, a Data Scientist creates a case management prediction that calculates the propensity of whether the case is likely to miss the SLA.

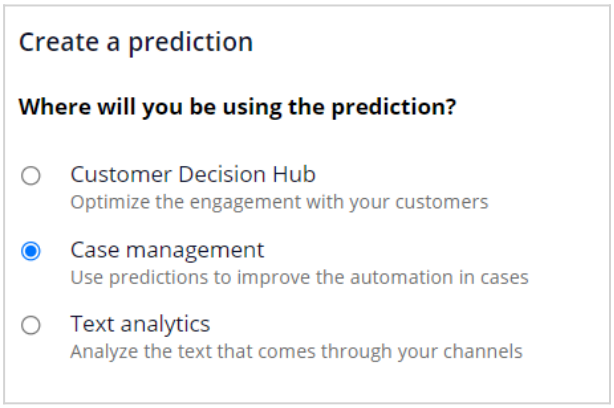

Process AI offers a wizard to create Missing Service-Level Agreement (SLA) predictions.

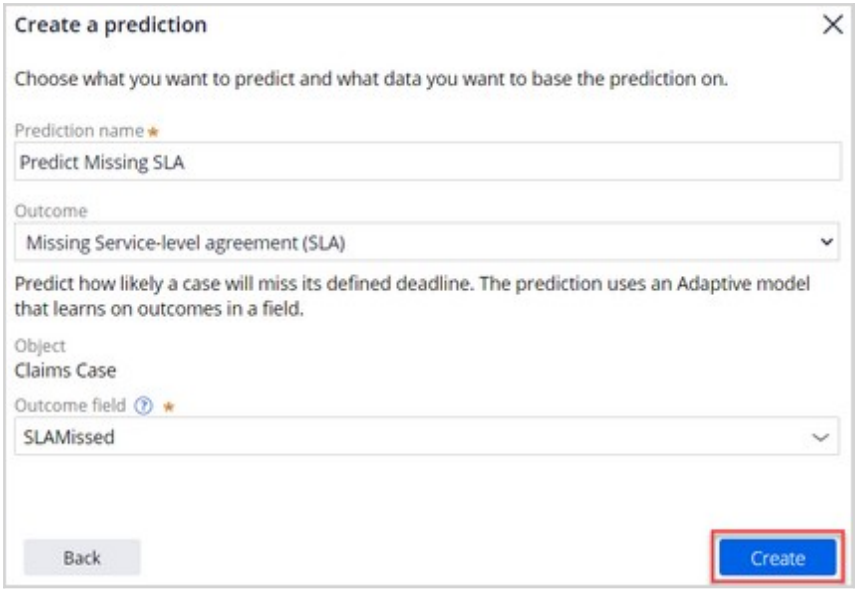

The Outcome field reflects the Boolean field that the Application Developer creates. This associates the prediction with the case type.

Next step for the Data Scientist is to add potential predictors. Best practice is to include many unrelated fields, including the claim properties. It is also important to exclude predictors that are irrelevant and do not have any predicting power, like **ChassisNo**, **CustomerID**, **CustomerPhoto**, **Name**, **PhoneNo**, **PolicyNo**, and **RegistrationNo**.

Add predictors

Click on the page and select fields

- Current page (InsuranceApp)
  - Page AccidentCategory_ea2f6
  - Page Customer
  - Page PredictMissingSLA

Name	Data type
AccidentDate	Date
Address	Text
Age	Integer
AgeOfCar	Double
ChassisNo	Text
ClaimedAmount	Double
ClaimFrequency	Integer
ClaimRatio	Double
ClaimType	Text
CustomerID	Identifier
CustomerPhoto	Text
DaySinceLastContacted	Integer
DOB	Date
DocumentsSubmitted	TrueFalse
DriverType	Text
EngineNo	Text
EntitledAmount	Double
Gender	Text
HospitalizationExpenses	Decimal
InjuryStatus	Text
Name	Text
NoClaimBonus	Integer
NoOfCarsInvolved	Integer
Occupation	Text
OdometerReading	Text
PedestrianInvolved	Text
PharmacyBills	Decimal
PhoneNo	Text
Placeofaccident	Text
PoliceReportAttached	Text
PolicyNo	Identifier
PreHospitalization	Decimal
RegistrationNo	Text
RelationToInsured	Text
Suspiciousclaim	Integer
Tenure	Double
TPAID	Text
TypeOfCar	Text
VehicleModel	Text
VehicleSpeed	Text

The adaptive model learns from previous cases and automatically activates predictors that perform above a threshold and deactivates predictors when their performance drops over time. The prediction is ready to be implemented in the Claims case case type by an application developer.

In the current configuration, the Decide complexity decision step categorizes claims as low or high complexity depending only on the claimed amount. As a result, claims that exceed 45000 are categorized as complex.

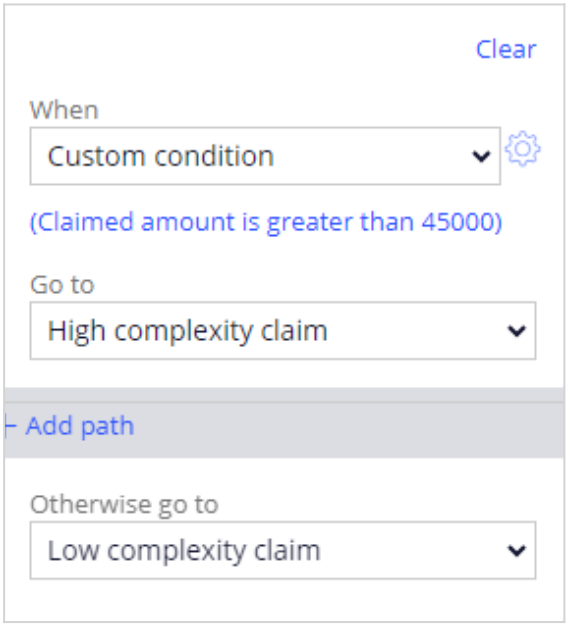

This condition requires an edit to meet the new business requirement that the routing decision is based on the propensity calculated by the Missing SLA prediction. To categorize a claim as a high complexity claim in the **Decide complexity** decision step, the propensity to miss the SLA needs to exceed a threshold. In this case: 0.4.

When a claims operator handles a claim, the case status changes to **Resolved-Completed** or **Resolved-Rejected**, and the outcome of the case maps to the alternative label (MissedSLA = False) for the prediction. When a complex claim misses the deadline, the outcome of the case maps to the target label (MissedSLA = True) for the prediction. The

model learns using this information and as a result, depending on the outcome, the missing SLA propensity for a similar case in the future increases or decreases.

A claim with a high propensity to miss SLA is immediately routed to an expert. The claim is routed to the regular workflow when the expert assesses the claim and does not consider it a complex case. This reassignment allows the adaptive model to learn from cases that are incorrectly routed to the expert.

An adaptive model is created for each primary and alternative stage in the case type. A decision request in a stage uses the model that is specific to that stage to calculate the propensity.

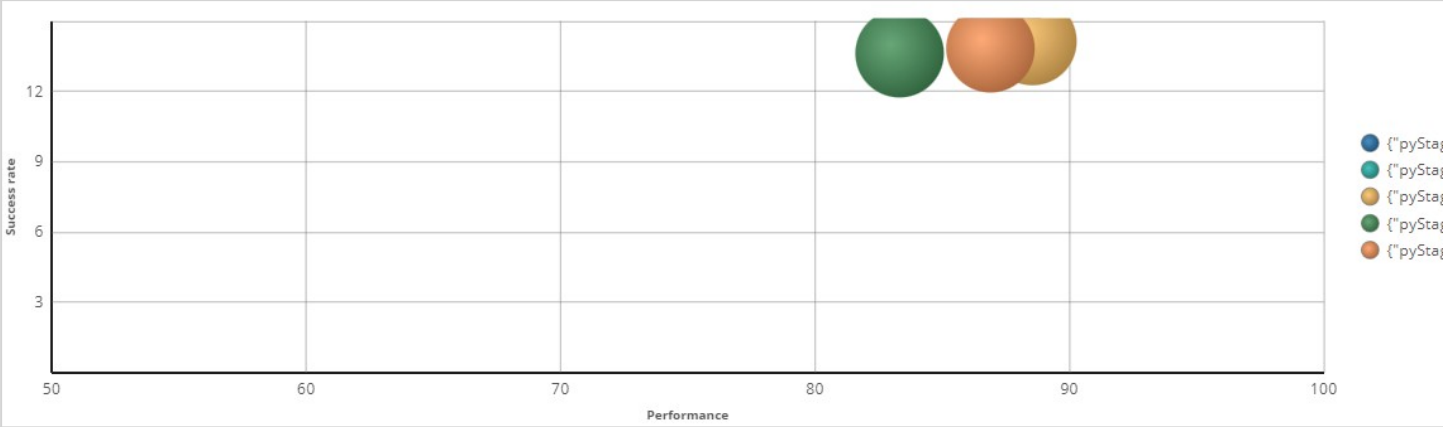

At the very beginning, the models have no predictive power. The models learn and self-optimize with every captured case outcome.

This demo has concluded. What did it show you?

- How to create a missing SLA prediction.
- How to implement a missing SLA prediction to improve efficiency.