

Make machine learning simple with Predictive Intelligence

What's in this Success Playbook

Machine learning (ML) holds significant promise for IT service management and customer service management, particularly for incident categorization, routing, and case assignment. In this Success Playbook, you'll learn how to:

- Use Predictive Intelligence (now available in Agent Workspace) to improve incident categorization, routing, case assignment, and solutions
- Demystify machine learning and get your service desk team in the comfort zone
- Create pilots with Predictive Intelligence
- Refine and scale your use of Predictive Intelligence for better incident and case management

Key takeaways

The most important things to know

- How to educate your teams on machine learning, both as a means to “demystify” the concept and spur innovative thinking
- Ways you can collaborate with your teams to take advantage of this thinking, and develop smart pilots focused on improving the performance of the incident and case management processes (rather than simply deploying a capability)
- How to use these pilots to build expertise and enthusiasm and to both scale and standardize the use of Predictive Intelligence

The payoff of getting this right

By using Predictive Intelligence effectively, you can significantly improve the accuracy and speed of incident categorization and case management, improving your mean time to resolve incidents and customer issues.

What you need to get started

Prerequisites

- Access to Predictive Intelligence
- A general understanding of [Predictive Intelligence](#)

When you should start this activity

Start on the initial steps in this Success Playbook as you begin to consider deploying Predictive Intelligence.

Step 1 – Educate your IT and customer service management organizations

Teach your service desk agents and customer service reps about machine learning.

KEY INSIGHTS

- Familiarize your service desk team with machine learning concepts and with how Predictive Intelligence works. When service desk agents and customer service reps understand both the concept and application of machine learning, they'll be more likely to identify opportunities to improve incident management and case assignment workflows.
- Don't focus on "selling" machine learning. Instead, work with your teams to modernize and align your incident management (ITSM) and case management (CSM and HR service delivery) processes using Predictive Intelligence—and identify where service desk agents, customer service reps, and HR staff can reallocate their time to higher-value tasks.

Machine learning can appear daunting to service desk agent, customer service reps, and HR staff, in worst-case scenarios, it can be perceived as a threat to service positions. Your use of Predictive Intelligence should start with an education campaign that focuses on the key insights listed for this stage.

Familiarize your staff with:

- Basic machine learning concepts
- How Predictive Intelligence works
- How machine learning technologies can improve their efficiency in ways that will let them shift from routine tasks to those with higher value

Position this education as both:

- A development opportunity for your teams – "Learn about the latest technologies."
- A change management component – "Machine learning will improve our team's performance."

Familiarize staff with basic machine learning concepts

Before you introduce Predictive Intelligence, create a primer on machine learning for your teams and cover it in your training—maybe as a team webinar, "lunch and learn," or workshop. Include, at a minimum, the following concepts:

- **A "plain language" definition of machine learning** – The machine learning built into the Now Platform is based on a simple premise: Predictive Intelligence uses "good" historical data to make predictions and to automate decisions based on those predictions.
- **A "plain language" definition of key terms** – Ensure your teams are familiar with these concepts:
 - **Supervised (and unsupervised) learning** – A machine learning program can be "trained" in supervised learning to make predictions or decisions based on a predefined data set, so that it makes similar predictions or decisions accurately when it's given new

data. In unsupervised learning, a machine learning program finds patterns or relationships in a new data set without prior training. Predictive Intelligence includes both types of machine learning capabilities.

- **Training** – Training refers to giving a machine learning program data that it can learn from by finding patterns that match input data to the answer it wants to predict.
- **Classification** – Classification is a kind of supervised learning that involves the machine learning program taking an input (like an incident) and assigning a label to it (for example, high, medium, or low severity).
- **How machine learning applies to IT incident management** – Effective prediction provides (a) support for major incident detection, by analyzing similarities between open incidents, (b) accurate incident categorization, enabling automated ticket assignment to the right assignment groups, and (c) recommended solutions based on past incidents.
- **How machine learning applies to customer case management** – Effective prediction provides (a) support for major incident detection, by analyzing similarities between open cases, (b) more accurate assignment of cases to the right groups, and (c) recommended solutions based on past cases.
- **How machine learning applies to HR service delivery** – Effective prediction provides (a) more accurate assignment of cases to the right CoEs/groups and (b) recommended solutions.
- **What machine learning enables for agents and reps** – When you use machine learning effectively, it:
 - Improves the productivity of staff because it reduces the time they waste on incorrect ticket routing and case assignment, and auto-generates recommended solutions
 - Reduces the time required to train staff on routine tasks, like incident categorization and case management
 - Improves triage, reducing outages and the time staff spend on firefighting activities
- **How it improves customer satisfaction and SLA compliance** – When you use machine learning effectively, it:
 - Improves speed of incident and case resolution, by reducing the time required to categorize, assign, and resolve incidents and cases
 - Improves customer satisfaction, through improved, early detection (and resolution) of major incidents, and reduced error rates in incorrect routing and assignment

Familiarize staff with how Predictive Intelligence works

When you provide your team basic groundwork in machine learning fundamentals, you prepare them to understand how Predictive Intelligence works. See Figure 1 for an overview that you can use for team education.

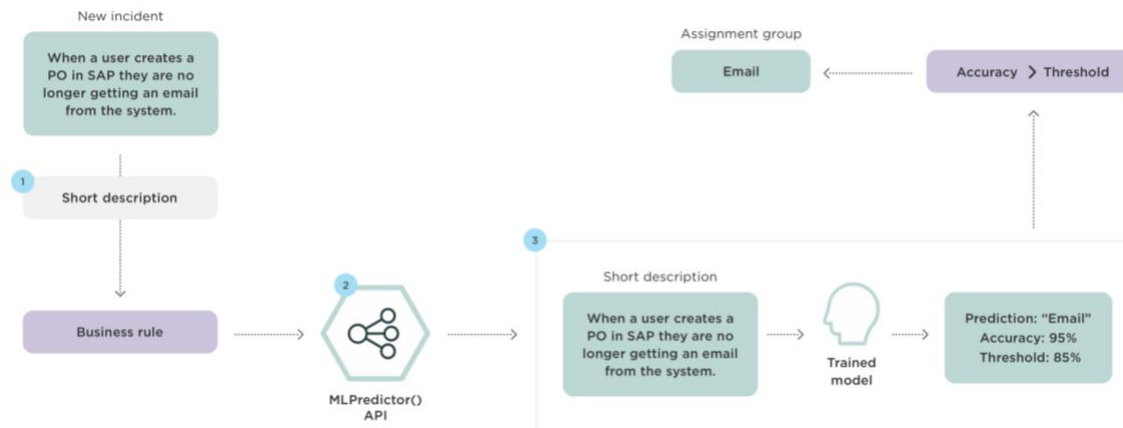


Figure 1: The Predictive Intelligence process for incident categorization

Predictive Intelligence begins with a business rule that triggers a call to the Prediction API for a new incident (see step 1 in the diagram).

The Prediction API takes an input—in this case, the short description for the incident (see Step 2), and runs it through a trained model to generate two outputs: a prediction about the assignment (step 3). If the confidence of the prediction is higher than a threshold value, then the field value for the assignment is automatically set (step 4).

You can base the input and output on solution definitions, such as incident assignments, that are already built into Predictive Intelligence, or they can be configured by your organization for specific outcomes. Any configuration should be based on a short statement or hypothesis that will drive the machine learning model.

Here's an example statement for incident assignment with a screen shot that shows the relevant fields from the Predictive Intelligence Solution Definition screen:

Statement – The short description field (input) for new incidents will be used to determine the incident assignment group.

The screenshot displays the 'ML Solution Definition - Incident Assignment' configuration page. Key elements include:

- Solution Template:** Incident Assignment Template
- Table:** Incident [Incident] (highlighted with a red box)
- Domain:** global
- Filter:** 0 records match condition. Below this, a list of conditions is shown, all of which must be met:
 - Active is false
 - Created on Last 12 months
 - State is one of New, In Progress, On Hold, Resolved (highlighted with a red box)
- Input Fields:** Short description (highlighted with a red box)
- Output Field:** Assignment group (highlighted with a red box)
- Training Frequency:** Run Once
- Label:** Incident Assignment
- Name:** ml_incident_assignment
- Active:** checkbox (unchecked)

Figure 2: Screenshot of an Predictive Intelligence solution definition with the relevant input, output, and condition fields highlighted

For additional demos you can use to support your team's education, visit our [Predictive Intelligence](#) page.

Familiarize service desk staff with the value of Predictive Intelligence

Before you introduce Predictive Intelligence, get your service desk team together and discuss the role that Predictive Intelligence should play in addressing the team's objectives. The purpose of this discussion is to solicit input on:

- How best to incorporate Predictive Intelligence within your current service desk processes and workflows
- How to incorporate Predictive Intelligence in a way that improves the team's ability to achieve its key performance indicators (KPIs)
- Opportunities for service desk staff to reallocate the time Predictive Intelligence saves them to higher-value activities

The intent of this exercise should not be to "sell" Predictive Intelligence. Instead, use this discussion to get your service desk agents to collaboratively plan how they can improve workflows and team performance using Predictive Intelligence. The best method for this is through a workflow mapping or process re-engineering workshop. With this approach, your team is more likely to walk away with a true understanding of Predictive Intelligence's potential value, including its possible process and performance improvement opportunities.

Step 2 – Develop a pilot to demonstrate Predictive Intelligence's capabilities

Start with an initial use case to define how you can apply machine learning.

KEY INSIGHTS

- Build and focus a pilot on a specific problem or use case so you can measure success through improved mean time to resolve/repair (MTTR), reductions in case and incident rerouting, or major incident detection.
- Start with out-of-the-box solutions, but include different approaches in your pilot to arrive at the optimal configuration.

The discussion or workshop you have with your service desk team (see the end of Step 1) should begin to surface parameters for a pilot that demonstrates Predictive Intelligence's capabilities—and validates its process and workflow improvement opportunities.

Planning the pilot

Focus your pilot on using Predictive Intelligence to address one or more of the following use cases: incident classification/categorization, use of a similarity framework to auto-generate recommended solutions, use of clustering to support major incident detection, and use of clustering to identify Knowledge Base gaps and opportunities. The capabilities to support these use cases are available under Predictive Intelligence as shown in Figure 3.

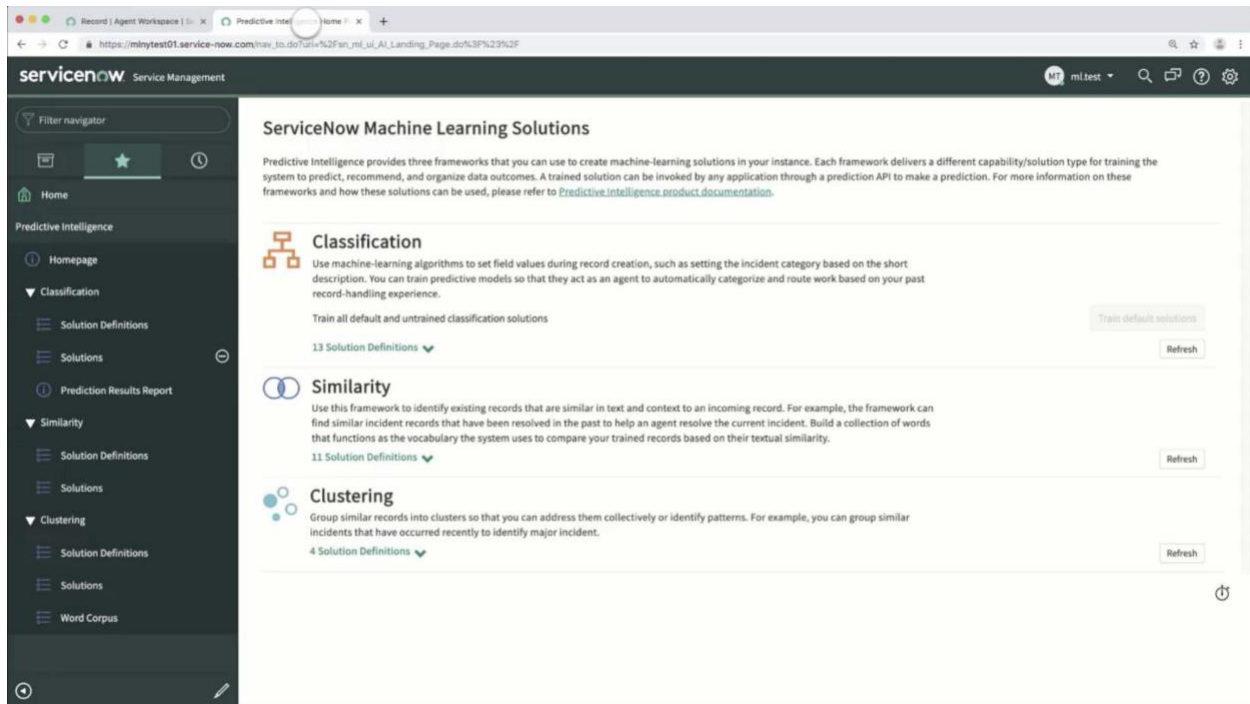


Figure 3: Screenshot of Predictive Intelligence solutions and the navigation menu

Use Case 1: Automatically categorize and automatically assign incidents (for IT staff) or cases (for customer service reps or HR staff). Make these the three critical elements of your pilot:

- **Explicitly define a use case or problem to be solved** – Consider limiting your pilot's scope to a specific category of incidents and cases, like a category within IT infrastructure or hardware or category of customer service case.
- **Explicitly identify the KPIs tied to your use case** – Try measuring success by one or more of the following KPIs:
 - **Mean time to resolve/repair (MTTR)** – By increasing the speed of ticket creation and assignment, incidents addressed through Predictive Intelligence should see improvements in MTTR relative to historical performance.
 - **Percentage of tickets or cases rerouted** – By increasing the accuracy of categorization and assignment, incidents and cases addressed through Predictive Intelligence should see less rerouting relative to historical performance.

In your pilot, you may want to include a comparison between historical workflows and the new workflows that use Predictive Intelligence. Alternatively, you can create a control group that isn't using Predictive Intelligence for categorization and routing.

- **Develop a good training model**¹ – The success of your pilot will be contingent on the data you select to build a training model. Good data ensures that the proper predictors “explain” the right categorization—for example, data from hardware-related incidents (e.g., type, age, etc.) that predict an appropriate incident category. The temptation may be to include more data where possible, but “more data” is not necessarily “good data,” from the standpoint of a training model. Instead, focus on three questions to ensure that the data you need addresses the problem or use case your team is trying to solve:

¹ For the purposes of this step, we'll refer to an example using Predictive Intelligence for ITSM incident management.

- What is the extent of data we have available? Get a clear picture of all data that you can use, including CMDB data, other database tables, and data from connected systems (including monitoring data).
- What data is not available that may be useful? You may be able to derive or simulate data that is not or cannot be recorded.
- What data can you omit and still address the problem? Excluding data is almost always easier than including data, and it often results in a better training model. Document any excluded data along with an explanation of why.

You will likely have assumptions about the data you include in your training model, such as why it reflects “good data” and how you anticipate it will help prediction. Record those assumptions so you can test them later as needed. For “good data,” focus on longitude—try to use at least a year’s worth to help prediction.

Use Case 2: Use a similarity framework to auto-generate recommended solutions for incidents (for IT staff) or cases (for customer service reps or HR staff). Make these the three critical elements of your pilot:

- **Explicitly define a use case or problem to be solved** – Consider limiting your pilot’s scope to a specific category of incidents and cases, like a category within IT infrastructure or hardware or category of customer service case.
- **Explicitly identify the KPIs tied to your use case** – Try measuring success by one or more of the following KPIs:
 - **Mean time to resolve/repair (MTTR)** – By auto-generating recommended solutions, IT service desk staff and/or customer service reps should be able to resolve incidents faster, resulting in improvements in MTTR relative to historical performance.
 - **Percentage of tickets or cases reopened** – By increasing the accuracy of solution recommendations, incidents and cases addressed through Predictive Intelligence should see fewer “reopens” relative to historical performance.

In your pilot, you may want to include a comparison between historical workflows and the new workflows that use Predictive Intelligence. Alternatively, you can create a control group that isn’t using Predictive Intelligence for solutioning.

- **Develop a good word corpus** – A similarity framework provides a vocabulary that Predictive Intelligence can use to compare new incidents and cases to past incidents and cases. To do this, you have to create a *word corpus* of text and context using the incident or asset table. Select record fields that contain the text or context you want in your word corpus. For example, if you’re trying to configure a solution to find similar incidents (such as “printer broken”), you may want to select fields like **short description**, **description**, **resolution notes**, and **close notes**. You then want to identify “similarity fields” that are likely to contain those words and phrases. Again, record your assumptions around the fields you select so you can test them later as needed – you’ll want to ensure that you have sufficient data in these fields for more accurate recommendations.

Use Case 3: Use clustering to improve detection times for major incidents (for IT staff). Note that this use case is unsupervised machine learning, meaning that it does not have labeled outputs (such as assignment groups or incident categories). Rather, the goal is to infer the natural structure or clusters within a dataset.

Make these the three critical elements of your pilot:

- **Explicitly define a use case or problem to be solved** – Consider limiting your pilot's scope to a specific category of incidents and cases, like a category (or set of categories) within IT infrastructure or hardware or category of customer service case.
- **Explicitly identify the KPIs tied to your use case** – Try measuring success by one or more of the following KPIs:
 - **Mean time to resolve/repair (MTTR) for major incidents** – By discovering patterns among open incidents, you should be able to more quickly detect and triage major incidents, relative to historical performance.

In your pilot, you may want to include a comparison between historical workflows and the new workflows that use Predictive Intelligence. Alternatively, you can create a control group that isn't using Predictive Intelligence.

- **Develop a good word corpus** – To discover patterns among incidents. To do this, you have to create a *word corpus* of text and context, using historical incident records (up to 300,000 records). You should select the **short description** field to guide Predictive Intelligence as it clusters incidents. If you have limited your categories for your pilot, ensure this is reflected by selecting **Group by**, and then **category** in your clustering solution. Record your assumptions around the categories you select. You'll want to test that you've developed a good word corpus to support better clustering.

Use Case 4: Use a similarity framework and clustering to improve the quality and effectiveness of your knowledge base. Make these the three critical elements of your pilot:

- **Explicitly define a use case or problem to be solved** – Consider limiting your pilot's scope to a specific knowledge base, if you have more than one. This use case requires configuration of the Knowledge Demand Insights feature.
- **Explicitly identify the KPIs tied to your use case** – Try measuring success by one or more of the following KPIs:
 - **Number of knowledge gaps resolved** – Knowledge Demand Insights will find knowledge gaps by comparing knowledge bases with a task type.
 - **Reduction in redundant knowledge articles** – Use of Predictive Intelligence to spot related articles can help identify opportunities for consolidation, especially as staff create new articles.
 - **Improvement in incident and/or case deflection** – By identifying and resolving knowledge gaps, you should be able to deflect incidents and/or cases where previous knowledge was insufficient.

In your pilot, you may want to include a comparison between knowledge bases (assuming you have more than one in place) using one as a control group so you can demonstrate the capability of Predictive Intelligence in the other.

- **Develop a good word corpus** – For your similarity framework or clustering solution, your word corpus should use your published Knowledge Base articles. Ideally, your corpus should include both the short description and article body fields, to ensure you can compare potentially related knowledge articles and identify gaps.

Building and activating the solution

Predictive Intelligence contains predefined solution definitions, or you can build your own solution based on mandatory fields. When starting incident categorization, ServiceNow®

recommends using predefined solution definitions to ensure you understand the process for developing and tuning your training model. See Figure 4.

The screenshot shows the 'ML Solution Definition - Incident Assignment' interface. Key highlighted fields include:

- Table:** Incident [Incident]
- Filter:** 0 records match condition. Conditions: Active is false, Created on Last 12 months, State is one of New, In Progress, On Hold, Resolved.
- Input Fields:** Short description
- Output Field:** Assignment group

Figure 4: Screenshot of Predictive Intelligence solution definition with relevant input, output, and condition fields highlighted

For **incident categorization**, here's how to complete the fields highlighted in Figure 3:

- **Table** – Select the table containing your training model records or data. Machine learning solutions are only trained on your own data, which never leaves ServiceNow data centers.
- **Filter** – Select the conditions you want to apply to the training model records. In order to train a solution, the filter must return at least one record. If your filter returns no records, update it until it returns records you can use for training. Predictive Intelligence provides a default filter when you select a solution template. In general, a good filter has these characteristics:
 - The training records are inactive and have task states that represent completing work within your standard processes (e.g., resolved or closed).
 - The training records contain only correct values for the target field. Filter out records with unreliable target field values.
 - The training records contain multiple examples of each target field value you want the solution to predict.
 - The training records include common variations of the input fields.
- **Fields** – Select the input fields you want the solution to use to generate a prediction. The system provides default input fields when you select a solution template. In general, good input fields have these characteristics:
 - The fields are available to users when they create records. Configure the form to display all input fields.
 - The field data type is string. The more information a field provides, the more often a solution can make a prediction, and the more often the predictions are accurate.
 - The field has a default value. The field should not have a blank value. For example, all default solution definitions use the short description field as the input field.

- **Output Fields** – Select the field with the value you want the predictive model to set. The system provides a default output field when you select a solution template. In general, a good output field has these characteristics:
 - The field is a choice field or string field with a finite set of possible values.
 - The field is assumed to have some causal connection to the input fields.

For **solution recommendations based on similarity frameworks**, here's how to complete the fields highlighted in Figure 5:

Figure 5: Screenshot of Predictive Intelligence similarity framework solution definition with relevant input, output, and condition fields highlighted

- **Table** – Select the table containing the records you want to compare to other similar records—in this case, your Incident[incident] table.
- **Filter** – Select the conditions you want to apply to the training model records. In order to train a solution, the filter must return at least one record. If your filter returns no records, update it until it returns records you can use for training. Predictive Intelligence provides a default filter when you select a solution template. In general, a good filter has the same characteristics as identified in the incident categorization use case:
 - The training records are inactive and have task states that represent completing work within your standard processes (e.g., resolved or closed).
 - The training records contain only correct values for the target field. Filter out records with unreliable target field values.
 - The training records contain multiple examples of each target field value you want the solution to predict.
 - The training records include common variations of the input fields.
- **Fields** – Select the input fields you want the solution to use from your word corpus to generate a prediction. The system provides default input fields when you select a solution template. In general, good input fields have these characteristics:

- The fields are available to users when they create records. Configure the form to display all input fields.
- The field data type is string. The more information a field provides, the more often a solution can make a prediction, and the more often the predictions are accurate.
- The field has a default value. The field should not have a blank value. For example, all default solution definitions use the short description field as the input field.

For **clustering in support of major incident detection**, here's how to complete the fields highlighted in Figure 6:

The screenshot displays the configuration interface for a Predictive Intelligence clustering solution. The 'Table' is set to 'Incident [incident]'. Under 'Fields', 'Short description' is selected from the available list. The 'Use Group By' checkbox is checked, and 'Category' is selected as the group by field. The 'Update Frequency' is set to 'Every 1 day' and the 'Training Frequency' is set to 'Every 30 days'. A section titled 'STEP 3: Configure cluster properties' shows the 'Minimum number of records per cluster' set to '2'. Other visible elements include a filter section indicating '65 records match condition' and a stopwords section with 'Default English Stopwords' selected.

Figure 6: Screenshot of Predictive Intelligence clustering solution definition with relevant input, output, and condition fields highlighted

- **Table** – Select the table containing the records you want to group into one or more clusters—in this case, your Incident[incident] table, as it contains incident records you want to group together for a major incident analysis.
- **Fields** – Select the input fields you want the solution to use from your word corpus to identify the records you want to include in your cluster. For major incident detection, **short description** is recommended. You can also add **Filters** to Cluster Input Fields.
- **Group by** – When you select a value from this list, the system groups records into one or more clusters based on your selection. For this use case, **Category** is recommended. This will then group your incidents into clusters by category, aiding in identifying major incidents.
- **Refresh and Recluster Frequency** – In these fields, select how often you want the system to group new and updated records into clusters (aka **update frequency**) and how often you want the system to discard results and recreate clusters from the beginning (aka **training frequency**).

- **Minimum Number of Records per Cluster** – Use this field to set the minimum number of records you want to allow in any cluster – the minimum is 2.

To use a **similarity framework** to generate **related articles as a user creates or views a knowledge article**, here's what to do (note that for this to work, your Knowledge Base must include at least 1000 articles):

- **Activate the Predictive Intelligence for Knowledge Management plugin** (com.snc.knowledge_ml).
- **Select the Knowledge Similar Articles solution definition from your ML solutions list.** The Solution Definition form will be set to default **field values** for knowledge articles. You can adjust values as needed, as discussed in our earlier use case for the similarity framework, but with the plugin activated, you should be able to proceed with default values. Click **update & retrain**. When the solution is complete, related Knowledge articles appear in the Knowledge results section on the Knowledge form and in the Related Articles section. The Knowledge form appears when you create an article. The Related Articles section appears on the article view page in the Knowledge Management Service Portal and Now Mobile applications.

To use **clustering** to **identify Knowledge Base gaps**, here's what you need to do. First you will need to configure one similarity type solution definition, and then one clustering type solution definition for each task type: incidents, customer service cases, or other tasks.

- **Activate the Predictive Intelligence for Knowledge Management plugin** (com.snc.knowledge_ml). Use Knowledge Management Guided Setup to configure the Knowledge Demand Insights feature. This will include setting assignment rules for feedback tasks to resolve knowledge gaps, so you'll want to have decided on your assignment process in advance.
- **Navigate to Knowledge Demand Insights, and then Solution Definitions.** In the Solution Definitions (ML view) list, search for and select the similarity solution definition for the task type.
 - For customer service cases, select **Demand Insights: Similar Cases and Knowledge (ml_sn_global_similar_cases_and_kbs)**.
 - For incidents, select **Demand Insights: Similar Incidents and Knowledge (ml_sn_global_similar_incidents_and_kbs)**.
 - For tasks other than customer service cases and incidents, click **New** to create another similarity solution definition.
- On the **Similarity Definition** form, **verify the default field values** for customer service cases or incidents, or fill in the values for a custom configuration. Submit (or update) and train your solution definition.
- Next, search for and select the **clustering solution definition** for the task type.
 - For customer service cases, select **Demand Insights: Case Clusters Need Knowledge (ml_sn_global_cases_need_knowledge_cluster)**.
 - For incidents, select **Demand Insights: Incident Clusters Need Knowledge (ml_sn_global_incidents_need_knowledge_cluster)**.
 - For tasks other than customer service cases and incidents, click **New** to create another clustering solution definition.

- On the Clustering Definition form, verify the default field values for customer service cases or incidents, or fill in the values for a custom configuration. Submit (or update) and train your solution definition. Use the demand insights dashboard for your tasks to analyze knowledge gaps.

Regardless of use case, your pilot should try several different ML solutions to arrive at the optimal configuration. Your training model should achieve a balance between the number of training model records and the accuracy of predicted outputs. See Figure 7. From the ML Solutions tab, agents can find Solution Definitions and Solution Statistics pages to review and tune machine learning results.

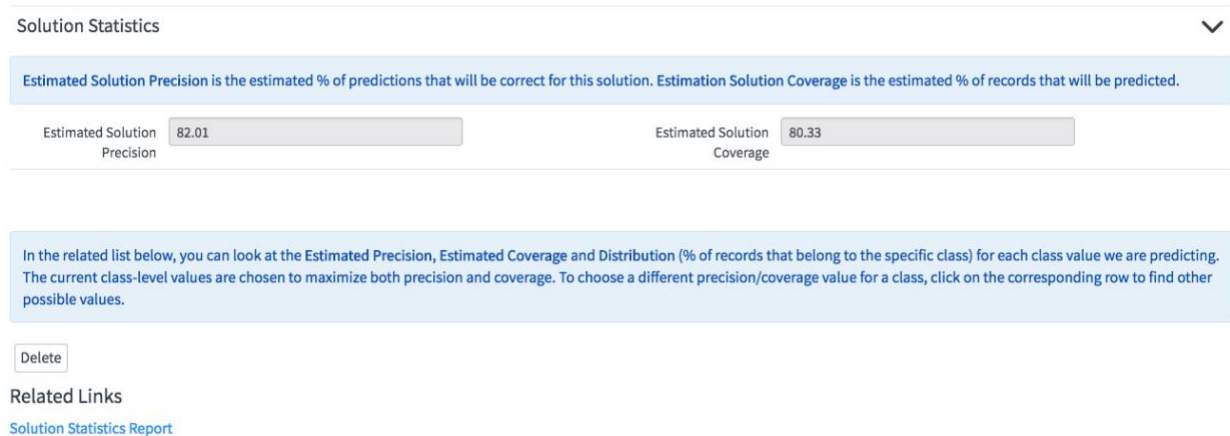


Figure 7: Screenshot of Predictive Intelligence solution statistics with estimated solution precision and estimated solution coverage

The key to tuning is finding the right balance of:

- **Estimated solution precision** – The estimated percentage of predictions that will be correct for your solution
- **Estimated solution coverage** – The estimated percentage of records that will be predicted

These two values are like two different dials on a radio that you can use to get to the station you want—you can generate more predictions (coverage) with less precision or generate greater precision to result in fewer predictions. Figure 8 shows a range of classes (such as a database) for a solution, with estimated precision and estimated coverage. For the database class, for example, incident categorization will cover 95% of the records with a 95% accuracy rate.

Related Links

[Solution Statistics Report](#)

	Class Confidence	Go to	Name	Search		1 to 6 of 6
Solution = ml_incident_categorization						
			Name	Estimated Precision	Estimated Coverage	Distribution
<input type="checkbox"/>			database	95	95	15
<input type="checkbox"/>			hardware	95	98	12
<input type="checkbox"/>			inquiry	50	47	28
<input type="checkbox"/>			network	95	94	14
<input type="checkbox"/>			request	95	100	14
<input type="checkbox"/>			software	95	100	14

Actions on selected rows... 1 to 6 of 6

Figure 8: Screenshot of Predictive Intelligence solution statistic for incident categorization, with the solution's Estimated Precision and Estimated Coverage across different classes (Note: The Distribution column refers to the percentage of records that belong to that class)

Once you find the right precision-to-coverage balance for your solution, you can train and activate a solution (or different solutions) to make predictions when new incidents or cases are created. Track your results over a set period of time to determine whether your solution is delivering enough correct predictions to improve the KPIs defined for your pilot.

EXPERT TIP

Before you move on to Stage 3, communicate the results of your pilot, including its effects on KPIs, to process owners, service desk staff (ideally those included in the workshop referenced in Stage 1), and senior stakeholders. Let this group know when a pilot was successful—it will help you build their support for moving Predictive Intelligence from pilot to regular use.

Step 3 – Build and refine the data you need to scale

Develop processes to revisit, review, and sharpen the data Predictive Intelligence uses.

KEY INSIGHTS

- Evaluate your pilot to improve your prediction accuracy and the number of predictions generated with some benchmarks that define success for the owner group.
- Continue to train a similarity framework to compare existing records to new similar records, so that you can reuse incident and case solutions.
- Increase your team's use of Predictive Intelligence to effectively embed the tool in your incident and case management processes.

Run a post-pilot evaluation to examine how you can improve your prediction accuracy and the number of predictions you generate. Completing this evaluation ensures that you scale your solution and incorporate it into the incident and case management workflows.

To improve prediction accuracy:

- **Eliminate stale data in your training model** – You can reduce the time frame in your training data by three months to eliminate stale incidents and cases. You should also identify and filter out categories and assignment groups from your training data that are no longer active.
- **Pick a target confidence value with lower coverage** – Your solution will make fewer predictions, but it'll make them with greater accuracy.
- **For categorization, review your solution to check why incidents are being reassigned or miscategorized** – For example, the solution bases its accuracy on where the incident is assigned when it's closed.

If a specific type of incident in your training data is (accurately) assigned to assignment group A, then reassigned to assignment group B for additional work, then closed, the solution will attempt to skip assignment group A and assign this incident type directly to assignment group B. This means that your data should reflect a process to reassign the incident back to assignment group A when it's closed (and your training data should be updated to reflect this).

To improve the number of predictions generated:

- **Pick a lower target confidence value with more coverage** – While the accuracy will decrease, your solution will generate more predictions.
- **Expand your training data** – Increase the time frame in your available data by three months. This may bring more stale data into your training model, so take steps to mitigate this, such as filtering out categories and assignment groups that are no longer active.
- **Check to see if something else is setting the output field data** – A prediction will only be made if this field is empty.

- **Identify whether a new term has entered your incident and case management process, such that it is not reflected in your training data** – For example, your organization may have deployed a new system called “Alpha” after you developed your training model. Because the term “Alpha” is not reflected in training model data, your solution will be unable to generate predictions for incidents or cases containing the term “Alpha.” Categorize these manually and assign them for one month to build training data—and rebuild your solution to accommodate them.

Note that you can use a similar approach to improve solution recommendations generated through a similarity framework or clustering.

Scaling your use of Predictive Intelligence is a matter of revisiting, refining, and adding to your training data over time—and reconfiguring your solution definition. Ensure you work with service desk staff and your incident and case management process owners to:

1. **Tune your input fields.**

Understanding how your organization uses input fields (like “short description”) is essential to ensuring that the solution reflects the data your organization collects for new incidents and cases.

2. **Test new input/output combinations based on hypotheses.**

Reconfigure your solution definitions on a regular cadence to understand causal connections between your training data and solution predictions and to improve the coverage and accuracy of your solution(s).

3. **Retrain your solution as service desk processes (or business processes) change.**

Any change to your incident and case management process (such as a reorganization of assignment groups) should prompt you to rebuild your solution(s) to ensure that auto-categorization and routing reflect the new process.

Over time, this will require your incident and case management process owners and select service desk staff to become fluent with Predictive Intelligence capabilities, so that the development and refinement of training models and solutions becomes embedded in incident management processes. Update job descriptions, process documentation, and training curricula to reflect this, so that using machine learning becomes a need-to-have rather than nice-to-have capability.

The takeaway

Machine learning may be daunting (or even threatening) to teams that have long been accustomed to legacy, manual processes. But the fundamentals behind machine learning are relatively straightforward. If your organization wants to take advantage of the capabilities of ServiceNow Predictive Intelligence, take the time to:

- Educate your teams on machine learning, both as a means to demystify the concept and spur innovative thinking
- Collaborate with your teams to take advantage of this thinking, and develop smart pilots focused on improving the performance of the incident management or case assignment process, or develop automated solution recommendations (rather than simply deploying a capability)
- Use these pilots to build expertise and enthusiasm and to both scale and standardize the use of Predictive Intelligence. Consider the development and use of dedicated champions for machine learning in your organization.

By following these steps, you can move machine learning from something perceived as advanced to something that's standard in process and workflow. Your teams will build the understanding that machine learning and automation are keys to better productivity—and the service desk has more opportunities to deliver value to the enterprise.