




From Bayes Theorem to Confusion Matrix

September 29, 2026

A confusion matrix supplies empirical counts of errors, while Bayes' theorem explains how those error rates and the class prevalence determine the probability that a positive prediction is actually correct. But there is a transformation from Bayes to CM and vice versa.

The study has 7 parts (Tutorial 199 of maXbox Starter Series):

1. Confusion Matrix
2. Where Bayes enters
3. Numerical Example (PoC)
4. Important Distinction
5. Explainer Video
6. Source Code Demo (Interactive maXbox Java Script)
7. Conclusion and Short History

The key distinction is:

- **False-positive rate:** "Among truly negative cases, how many were incorrectly flagged?"
- **Precision / positive predictive value:** "Among all cases flagged positive, how many are real truly positive?"

These are not the same quantity.

Confusion Matrix

Fundamental Metrics and Definitions

To understand this connection, it is critical to distinguish between metrics that evaluate negative cases versus metrics that evaluate flagged cases:

- **False-Positive Rate (FPR):** Answers the question, *"Among truly negative cases, how many were incorrectly flagged?"*
- **Precision / Positive Predictive Value (PPV):** Answers the question, *"Among all cases flagged positive, how many are actually true positives?"*
- **Recall / True Positive Rate (TPR):** Answers the question, *"Among all actual positive cases, how many were correctly identified?"*
- **Prevalence:** The baseline proportion of positive cases within the total population.

For a binary classifier:

	Predicted positive	Predicted negative
Actually positive	TP	FN
Actually negative	FP (False Positive)	TN

The standard metrics are true positive rate, false positive rate and precision:

$$\text{TPR} = \text{Recall} = \frac{TP}{TP + FN}$$

$$\text{FPR} = \frac{FP}{FP + TN}$$

$$\text{Precision} = \text{PPV} = \frac{TP}{TP + FP}$$

Precision=PPV=TP/(FP+TP).

$$\text{TPR} = \text{Recall} = \frac{TP}{TP + FN}$$

$$\text{FPR} = \frac{FP}{FP + TN}$$

$$\text{Precision} = \text{PPV} = \frac{TP}{TP + FP}$$

A confusion matrix records these counts directly; for scikit-learn's binary convention, false positives are actual class 0 predicted as class 1. [[scikit-learn](#)]

Where Bayes enters

Let:

- D : the object or condition is genuinely positive.
- $+$: the classifier predicts positive.
- $P(D)$: prevalence, the fraction of truly positive cases.
- $P(+ | D)$: true-positive rate, or recall.
- $P(+ | \neg D)$: false-positive rate.

The question “If the classifier says positive, how likely is it to be truly positive?” is:

$$P(D | +).$$

Bayes' theorem gives:

$$P(D | +) = \frac{P(+ | D)P(D)}{P(+ | D)P(D) + P(+ | \neg D)P(\neg D)}.$$

Substituting the confusion-matrix quantities:

$$\text{Precision} = \frac{\text{Recall} \cdot \text{Prevalence}}{\text{Recall} \cdot \text{Prevalence} + \text{FPR} \cdot (1 - \text{Prevalence})}$$

$$\boxed{\text{Precision} = \frac{\text{Recall} \cdot \text{Prevalence}}{\text{Recall} \cdot \text{Prevalence} + \text{FPR} \cdot (1 - \text{Prevalence})}}$$

This is the direct connection. Precision is the Bayesian posterior probability of a real positive after observing a positive prediction.

Numerical example

Suppose an object detector has:

- Recall = **0.90**
- False-positive rate = **0.05**
- Positive prevalence = **0.01**, meaning only 1% of image regions really contain the object.

Then for Bayes:

$$P(D \mid +) = \frac{0.90 \cdot 0.01}{0.90 \cdot 0.01 + 0.05 \cdot 0.99}$$

$$= \frac{0.009}{0.0585} \approx 0.154.$$

$$P(D \mid +) = \frac{0.90 \cdot 0.01}{0.90 \cdot 0.01 + 0.05 \cdot 0.99}$$

$$= \frac{0.009}{0.0585} \approx 0.154.$$

So the precision is only about 15.4%, despite a respectable 90% recall and only 5% false-positive rate.

For 10,000 tested regions:

Category	Count
True positives	$100 * 0.90 = 90$
False negatives	$100 - 90 = 10$
True negatives	$9,900 * 0.95 = 9,405$
False positives	$9,900 * 0.05 = 495$

Confusion Matrix	Predicted positive	Predicted negative	Sum
Actually positive	90 (TP)	10 (FN)	100
Actually negative	495 (FP)	9405 (TN)	9900
Sum	585	9415	10000

Therefore:

$$\text{Precision} = \frac{90}{90 + 495} = 15.4\%.$$

$$\text{Precision} = 90 / (495 + 90) = 15.4\%.$$

The detector generates many more false alarms than correct detections because negative cases vastly outnumber positive ones. This prevalence effect is why precision is sensitive to class imbalance, whereas FPR is normalized only by the number of actual negatives.[\[scikit-learn\]](#)[\[cloudspress\]](#)

Why false positives matter so much

A small FPR can produce many false positives when the negative population is large.

For an object detector searching for a rare object—say a railway signal in a large aerial image—even an FPR of 0.1% may cause numerous false detections if millions of candidate regions are examined.

That creates two different operational questions:

- **FPR:** How often does the detector falsely flag a truly negative location?
- **Precision:** When the detector flags a location, how often is the alert correct?

Advertisement



Japanische Köche waschen nie ihr Schneidebrett – sie...

TrendSpotOnline

[Privacy Settings](#)

For an operator reviewing alerts, precision is usually the more intuitive measure. For comparing discrimination across thresholds and datasets, ROC analysis uses TPR versus FPR. For rare-object detection, precision–recall curves are often more informative because they show the effect of prevalence and false discoveries.[\[cloudspress\]](#)[\[lucdemortier.github\]](#)

Relation to your correlated-object model

Suppose a local detector proposes a weak fence:

$$P(\text{Fence} \mid \text{image patch}) = 0.48.$$

A nearby car is then used as contextual evidence. The contextual model should estimate something like:

$$P(\text{Fence} \mid \text{local image evidence}, \text{nearby car spatial relation}).$$

After applying the context, perhaps the score becomes 0.72. That value is a model score or posterior-like probability, not automatically a measured precision.

To verify whether 0.72 is meaningful, evaluate the final context-aware detector on labelled data and build a new confusion matrix. The contextual algorithm may:

- convert some previous false negatives into true positives;
- convert some previous false positives into true positives;
- or create new false positives if the correlation rule is too strong.

Thus, you need separate confusion matrices for:

1. The local detector alone.
2. The detector after correlated-object rescoreing.

Advertisement



**Japanische Köche waschen
nie ihr Schneidebrett – sie...**

TrendSpotOnline

[Privacy Settings](#)

Then compare:

$$\Delta TP, \quad \Delta FP, \quad \Delta FN, \quad \Delta TN,$$

$\Delta TP, \Delta FP, \Delta FN, \Delta TN,$

as well as precision, recall, FPR, and calibration.

Important distinction

A detector's internal Bayesian or softmax probability and its empirical confusion-matrix precision are different:

$$P(\text{true object} \mid \text{image and context})$$

$$P(\text{true object} \mid \text{image and context})$$

is the model's posterior estimate, while

$$\frac{TP}{TP + FP}$$

is the observed fraction of correct positive predictions on a particular labelled evaluation set.

They coincide only approximately if:


- the probability is well calibrated;
- the evaluation data has the same prevalence as the deployment data;
- the threshold and object-matching rules are identical;
- the test data represents real operating conditions.

Conclusion

So, in one straight sentence:

The confusion matrix measures false positives directly; Bayes' theorem converts the false-positive rate, recall, and prevalence into the **posterior probability** that a positive prediction is genuine—namely called **precision** of confusion matrix!

- At 0.01 prevalence, even 90% sensitivity + 95% specificity yields only ~15% precision.
- Raising specificity (lowering FPR) is far more effective than raising sensitivity when the positive class is rare.
- This is the Bayesian reason why threshold tuning and false-positive minimization matter in imbalanced detection.



Infographic based on "From Bayes Theorem to Confusion Matrix"

The infographic visually maps the mathematical connection between Bayes' Theorem and the confusion matrix, illustrating how recall, false positive rate (FPR), and class prevalence combine to calculate precision. It highlights the numerical example demonstrating how a low class prevalence (1%) leads to a precision of 15.4% despite a 90% recall and a low 5% FPR.

Max Kleiner, 29/09/2026

Advertisement



**Japanische Köche waschen
nie ihr Schneidebrett – sie...**

TrendSpotOnline

[Privacy Settings](#)

Explainer Video

The video will synthesize the key mathematical concepts, error metrics, and probability calculations.



The Precision Paradox Bayes ConfusionMatrix

Max Kleiner



Watch on

Bridge the gap between Bayes and Confusion Matrix

Source Code Demo

Here is a self-

contained **maXbox program** that demonstrates how **class** imbalance destroys precision, using the Bayesian relationship we discussed. It computes precision from sensitivity, specificity, **and** prevalence, **then** simulates a confusion matrix at different prevalence levels so you can see the effect numerically.

https://sourceforge.net/projects/maxbox5/files/examples/1499_bayes_confusionmatrix1.txt/download

It then simulates a population of 10,000 at five prevalence levels — from balanced (50%) down to 0.01% — **and** prints the resulting TP, FP, TN, FN counts.

The **final** section shows that at 1% prevalence, improving specificity from 0.95 to 0.9999 raises precision from ~15% to ~90%, while improving sensitivity from 0.90 to 0.999 barely moves precision. This is the practical consequence of the Bayes denominator: when negatives vastly outnumber positives, the false-positive term dominates, and only reducing FPR helps.

Interactive Form Demo App

maXbox5 - Bayes Theorem to C...

Total Population: 10000 Recall/TPR (0..1): 0.90
 Prevalence (0..1): 0.01 FPR (0..1): 0.05

Calculate Confusion Matrix Bayesian Precision

```

=== BAYES THEOREM TO CONFUSION MATRIX
DEMO ===
Input Population: 10000.000000000000
Class Prevalence P(A): 1.000000000000%
Recall (Sensitivity / TPR):
90.000000000000%
False Positive Rate (FPR):
5.000000000000%

--- CONFUSION MATRIX COUNTS ---
Actual Positive Count: 100.000000000000
Actual Negative Count: 9900.000000000000
  True Positives (TP) : 90.000000000000
  False Negatives (FN) : 10.000000000000
  False Positives (FP) :
495.000000000000
  True Negatives (TN) :
9405.000000000000
  Total Predicted Positive:

```

https://sourceforge.net/projects/maxbox5/files/examples/1499_bayes_confusion_matrix_demoform_nl.pas/download


Script Highlights:

- **Bayesian Precision & Confusion Matrix Modeling:** Translates the mathematical bridge connecting Recall, False Positive Rate (FPR), and Class Prevalence into empirical confusion matrix counts (\$TP\$, \$FP\$, \$TN\$, \$FN\$)¹².
- **Interactive maXbox GUI:** Provides editable controls (TEdit) for total population, class prevalence (\$P(A)\$), recall (\$TPR\$), and false positive rate (\$FPR\$), outputting detailed step-by-step calculations to a structured log view³⁴.
- **The Precision Paradox Visualized:** Demonstrates why a 1% class prevalence yields a precision of only 15.4% (\$90 / (90 + 495)\$) despite a 90% recall and a low 5% false positive rate³⁴.

Interactive Form Demo App V2

New Features in Version 2:

- **Embedded TImage Canvas Plotter:** Plots a real-time Precision vs. False Positive Rate (FPR) curve directly on the right-hand panel of the maXbox form [15].
- **Dynamic Operating Point Callout:** Highlights the active FPR and Precision (e.g., FPR = 5.0%, Precision = 15.4%) with a red coordinate marker and data label [15].
- **Prevalence Comparison Curve:** Overlay a reference curve for a higher prevalence (5.0%) alongside the active setting (1.0%), visually illustrating how rare target classes cause precision to drop steeply as FPR increases [15].




https://sourceforge.net/projects/maxbox5/files/examples/1499_bayes_confusion_matrix_demoform_n12.pas/download

Interactive Form Demo App V2.1

What was added in Version 2.1:

- **Real-Time TrackBar Slider (TrkFPR):** Adds a dynamic slider allowing you to sweep the False Positive Rate continuously from 0.0% to 20.0% (with 0.1% step resolution).
- **Interactive Axis Crosshairs & Moving Dot:** As you drag the slider, red crosshairs trace from the FPR and Precision axes to the red operating point, updating its coordinates on the plot in real time.

- **Synchronized Event Handlers:** Editing either the numerical FPR text box or dragging the trackbar automatically syncs both controls and recalculates the confusion matrix and plot without lag.




https://sourceforge.net/projects/maxbox5/files/examples/1499_bayes_confusion_matrix_demoform_n121.pas/download

Web Solution

This interactive web application translates the mathematical bridge between Bayes' Theorem and the Confusion Matrix into a real-time web calculator

<https://raw.githubusercontent.com/maxkleiner/maXbox5/main/assets/bayesmatrix.htm>



Bayes' Theorem to Confusion Matrix & Precision-FPR Plotter of maXbox5

Short History

Where it comes from

The confusion matrix is a 2x2 **contingency table** of predicted versus actual class membership. Each cell is a count, and every metric derived from it (sensitivity, specificity, precision, NPV) is a **conditional probability** estimated from those counts. Bayes' theorem is the general rule for inverting conditional probabilities, so the link is structural, not an external addition. [wikipedia](#)

The key identity is:

$$\text{Precision (PPV)} = P(\text{actual} + | \text{predicted} +) = \frac{P(+ | D+) P(D+)}{P(+ | D+) P(D+) + P(+ | D-) P(D-)}$$

which, after substituting the confusion-matrix quantities, becomes:

$$\text{PPV} = \frac{\text{Sensitivity} \times \text{Prevalence}}{\text{Sensitivity} \times \text{Prevalence} + (1 - \text{Specificity}) \times (1 - \text{Prevalence})}$$

This is exactly Bayes' theorem applied to the four cells of the confusion matrix.

Precision = PPV

Advertisements

Occasionally, some of your visitors may see an advertisement here, as well as a [Privacy & Cookies banner](#) at the bottom of the page. You can hide ads completely by upgrading to one of our paid plans.

UPGRADE NOW

DISMISS MESSAGE