Chapter 7

Confidence, Trustworthiness, and Composites

Collaboration can be essential to manage complex projects. One example is building a house. You then need the expertise of carpenters, plumbers, and electricians. Each profession brings unique skills to the table.

Similarly, different types of Tsetlin machines can have distinct capabilities. In this chapter, you learn how Tsetlin machines can team up, allowing them to achieve more than they could on their own.

The effectiveness of a team relies on recognising each member's strengths and limitations. Appreciating where your expertise stops and where your coworkers' expertise begins is crucial for effective collaboration. We first explore how Tsetlin machines can assess their competence in Section 7.1. Using the vote count from Chapter 1, you learn to measure how confident a Tsetlin machine is when it makes its decisions.

It is possible to be highly confident and still perform poorly. To be trustworthy, confidence must be in line with one's capabilities. Therefore, Section 7.1 also covers how to evaluate trustworthiness.

Next, in Section 7.2, you discover how to build a team of Tsetlin machines with different skills. By assessing each Tsetlin machine's confidence, you can lean on the confident ones when making decisions. The result is a *Tsetlin machine composite* – a construction where multiple Tsetlin machines join forces. You can think of it as a composite material, such as epoxy, which reinforces resin with fibres, making it strong, lightweight, and durable.

Figure 7.1: Visiting your family doctor.

7.1 Does a Tsetlin Machine Know When it Does Not Know?

Imagine that you have a health problem. To get help, you visit your family doctor (Figure 7.1). Your doctor will ideally have broad expertise but obviously cannot be an expert in all areas of medicine. Still, you can trust your doctor if they have one crucial ability: *They must know when they do not know.* If your doctor is uncertain about your condition, they can refer you to a suitable specialist for follow-up.

#	**Lectures**	**Textbook**	**Assignments**	***Pass***	**R1**	**R2**	**R3**	**v**
1.	yes	yes	yes	*yes*	•	•	·	+2
2.	no	yes	no	*no*	·	•	•	+0
3.	yes	yes	yes	*yes*	•	•	·	+2
4.	no	no	no	*no*	·	·	•	−1
5.	yes	yes	no	*yes*	•	•	•	+1
6.	no	no	no	*no*	·	·	•	−1

Table 7.1: Six students, their work habits, and exam outcome.

Similarly, if you receive medical advice from a Tsetlin machine, you would like to know when it is uncertain about its answers. You need a way to measure the confidence of your Tsetlin machine.

Confidence

Measuring confidence relies on the *vote sum* from Chapter 1. To demonstrate this, consider the task of predicting which of six students will pass or fail a course. Table 7.1 provides information on the habits of the students, with the columns **Lectures**, **Textbook**, and **Assignments** answering the questions:

- Did the student follow the lectures of the course?
- Did the student complete the textbook?
- Did the student solve the assignments?

Assume that you have a Tsetlin machine that uses this information to predict the pass/fail result with the help of three rules:

R1 **if** *Lectures* **then** *Pass*,

R2 **if** *Textbook* **then** *Pass*,

R3 **if not** *Assignments* **then** *Fail*.

The rules make the prediction by voting for Pass or Fail when their condition matches the information on the student. The bullets • in table columns **R1**, **R2**, and **R3** show the matches.

Recall how adding the Pass votes and subtracting the Fail votes determines the vote sum in column **v**. Figure 7.2 illustrates the relationship between the sum of votes and confidence, as explained below.

Confidence in Pass. If multiple rules vote for Pass and few against, you get a high vote sum. With many agreeing rules, we say that the Tsetlin machine is *confident.*

Uncertainty. Oppositely, if few rules vote or if they disagree, you get a vote sum closer to zero. In that case, there is no convincing consensus among the rules. We say that the Tsetlin machine is *uncertain.*

Confidence in Fail. The above reasoning also applies to Fail. However, the magnitude of the *negative* vote sum then assesses confidence.

Confidence Calculation. In general, it is the absolute value of the vote sum that measures how confident a Tsetlin machine is:

$$\text{Confidence} := \textbf{abs}(\text{Vote Sum}). \tag{7.1}$$

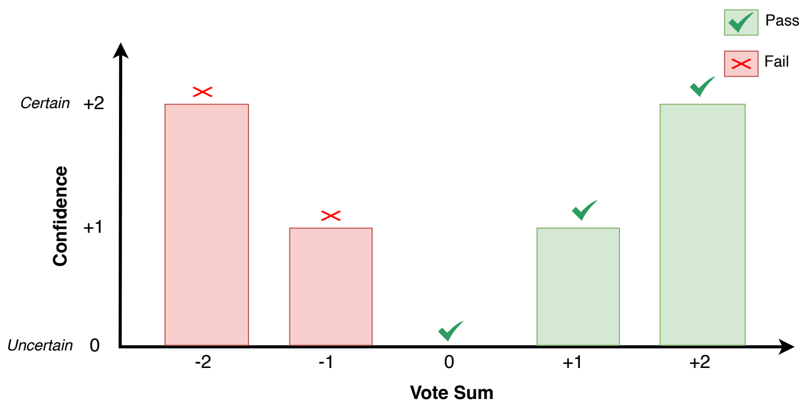

Figure 7.2: Relationship between vote sum and confidence.

Example of Confident Classification. Observe first how student #1 in Table 7.1 receives two Pass votes, one each from rules R1 and R2. It does not receive any Fail votes. Hence, the sum of votes is +2. The rules are all agreeing on the outcome of the majority vote, and the Tsetlin machine can be confident in Pass.

Example of Uncertain Classification. Now consider students #2. Rule R1 is silent, rule R2 is voting for Pass, and rule R3 is voting for Fail. The sum of votes is now +0. Here, we have disagreement, giving a tie. Accordingly, the Tsetlin machine is uncertain about whether student #2 will pass or fail.

Example of Confidence Ranking. Ranking the students from highest to lowest confidence, we get student #1, #3, #4, #5, #6, and #2. Students #1 and #3 achieve the highest ranking because the confidence is +2. Students #4, #5, and #6, in turn, obtain confidence +1. Finally, student #2 is classified with confidence +0.

Multi-Class Confidence. A multi-class Tsetlin machine merges one standard Tsetlin machine per class. Each produces a sum of votes for its class, and the class with the highest sum is the predicted class. Accordingly, the highest vote sum determines the confidence in the prediction:

$$\text{Multi-Class Confidence} := \mathbf{max}(\text{Vote Sums}). \tag{7.2}$$

Trustworthiness of Confidence

High confidence does not necessarily equate to good performance. As illustrated in Table 7.2, someone can be overly confident and still perform badly. On the other hand, being too timid can also lead to poor results

	Performance	
Confidence	Low	High
Low	*Justifiably Uncertain*	*Underconfident*
High	*Overconfident*	*Justifiably Confident*

Table 7.2: Trustworthiness of confidence.

due to lack of action. Confidence must be in line with one's abilities to be trustworthy.

Trustworthiness. For a classifier, you can measure the trustworthiness by comparing the actual class with the class assigned by the classifier. If high confidence leads to high classification accuracy, and vice-versa, confidence is faithful to performance.

Example of Trustworthiness Analysis. Table 7.3 performs this analysis for the Pass/Fail prediction task. First, contrast the actual class in column **Pass** of Table 7.1 with the vote sum in column **v**. Comparing the classifier's prediction with the actual class for the six students reveals the following:

- Student #1 and #3 enjoy a confidence of +2, and the classifier gives them the correct class. So, at this confidence level, the accuracy is 100%. Two out of two students get the correct class.
- Adding students #4, #5, and #6 to the group by decreasing the confidence level from +2 to +1 does not affect the accuracy of the classifier, which remains at 100%.
- For confidence level +0, the accuracy decreases to 83%. The Tsetlin machine correctly classifies five out of six students at this confidence level, misclassifying student #2.

In other words, acting with less confidence leads to more mistakes.

Image Classification Example. Figure 7.3 covers the confidence of a multi-class Tsetlin machine that classifies colour images into 100 distinct

Confidence	$\geq +0$	$\geq +1$	$\geq +2$
Correct	5 out of 6	5 out of 5	2 out of 2
Accuracy	83%	100%	100%

Table 7.3: The accuracy of the Tsetlin machine at different confidence levels.

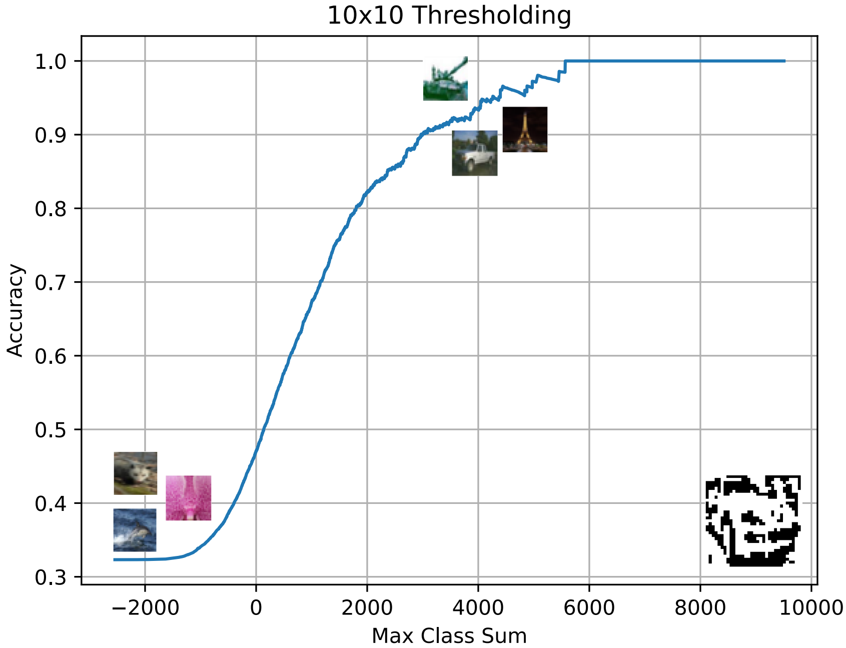

Figure 7.3: Trustworthiness of color image classification confidence.

classes.[1] This Tsetlin machine utilises a process known as pixel colour thresholding to convert the input into Boolean values. Each pixel is assigned the value True or False based on its colour and the colours of the neighbouring pixels. The image in the lower right of the figure shows an example output of this process.

The x-axis of the plot ranks ten thousand color images by multi-class confidence, ranging from −2 500 to 10 000. The y-axis then shows the accuracy of the multi-class Tsetlin machine from the x-axis confidence level and upward. Accordingly, this figure is analogous to Table 7.3 from the Pass/Fail prediction example. The difference is that we use multi-class confidence, which means that confidence increases with the vote sum of the

[1]We will return to image understanding in detail in Chapter 9.

class with the highest vote sum. Oppositely, confidence drops when the highest vote sum falls. Notice how higher confidence corresponds to higher accuracy. As confidence increases, the accuracy grows from approximately 34% to 100%, signifying trustworthiness.

Let us now look at the six colour images that annotate the plot. The Tsetlin machine is not sure about the three on the left. However, it is confident in the class assigned to the three on the right. As seen, this particular Tsetlin machine struggles with images mainly composed of colour textures, like the pink flower. On the other hand, it is highly confident when classifying larger objects, like a tank or the Eiffel Tower.

At this point, you know how to measure the confidence of a Tsetlin machine and how to assess its reliability, the foundation for building Tsetlin machine composites.

7.2 Building Tsetlin Machine Composites

Different Tsetlin machines can have different skill sets. To illustrate this, let us return to the example of predicting breast cancer recurrence from Chapter 1. Recall how the system must determine whether breast cancer will return after treatment (recurrence) by examining medical information on the patient. Here are two ways in which Tsetlin machines can obtain unique abilities.

Different Populations. Two hospitals may have trained their own Tsetlin machines to predict breast cancer recurrence. Differences in patient demographics, treatment approaches, and facilities will give the two Tsetlin machines unique perspectives.

Different Modalities. Tsetlin machines can also have different views on the patient. One Tsetlin machine may operate on blood samples, another on mammograms, and a third on the narrative written by the doctor. Each source of information offers valuable information on breast cancer recurrence.

Example Tsetlin Machines. Imagine that you have two Tsetlin machines. One has only seen patients with breast cancer recurrence, and the other has only seen non-recurrence patients.

- **Recurrence Tsetlin Machine.** This Tsetlin machine has only trained on recurrence patients, obtaining the two rules:

  R1 **if** *Deg-malign 3* **and not** *Menopause lt40* **then** *Recurrence*,

  R2 **if** *Deg-malign 3* **and not** *Menopause lt40* **then** *Recurrence.*

- **Non-Recurrence Tsetlin Machine.** This Tsetlin machine can only recognize non-recurrence patients using the following rule:

  R3 **if** *Inv-nodes 0-2* **then** *Non-Recurrence.*

You are now ready to investigate the final question: How can various Tsetlin machines work together to leverage their combined skill sets?

Tsetlin Machine Composites

Figure 7.4 shows the procedure for building a Tsetlin machine composite. We use the two example Tsetlin machines to illustrate the procedure. Let us look at each step in detail.

Tsetlin Machine Collaborators. Multiple Tsetlin machines participate in the composite. Within the composite, each Tsetlin machine operates independently in the standard way. However, instead of outputting a class, the Tsetlin machines output their vote sums. This means that any type of Tsetlin machine can be integrated into the composite in a plug-and-play manner.

Normalization. Different kinds of Tsetlin machines may operate on different data, with a varying number of rules, and with various vote margins. Accordingly, their confidence levels may become incompatible. The normalisation step aligns confidence across Tsetlin machines. First, each Tsetlin machine identifies the highest and lowest vote sums it obtains from the data. The difference between the upper and lower values becomes the normalisation factor for that Tsetlin machine. The Tsetlin machine normalises its vote sum by dividing with the normalisation factor:

$$\text{Normalized Vote Sum} := \frac{\text{Vote Sum}}{\text{Highest Vote Sum} - \text{Lowest Vote Sum}}. \tag{7.3}$$

For example, the highest vote sum of the Recurrence Tsetlin Machine is +2 while the lowest sum is +0 for the data in Table 1.9 of Chapter 1. Therefore, its normalisation factor is $2 - 0 = 2$.

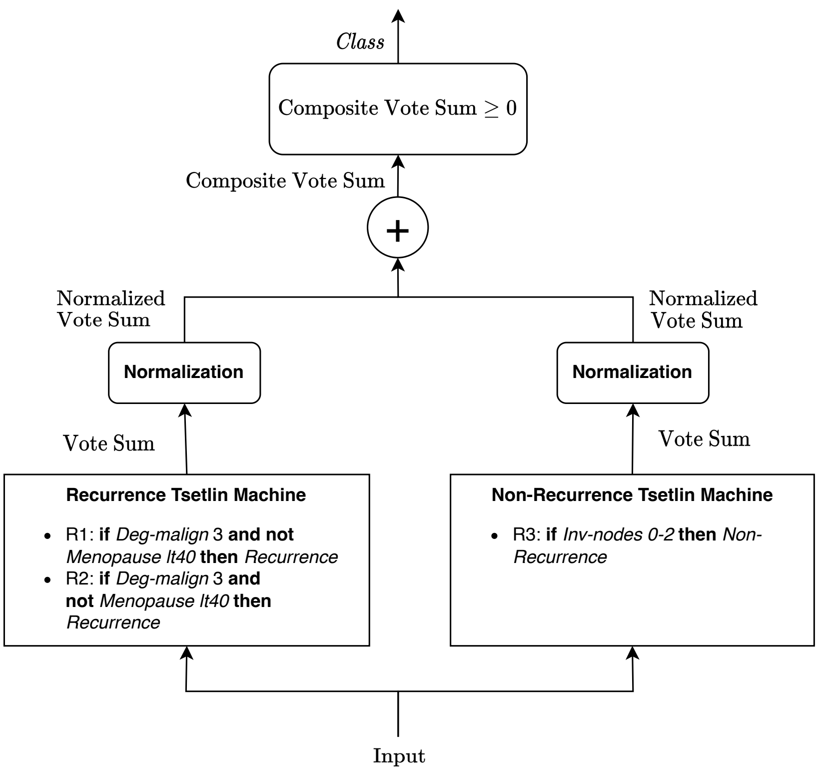

Figure 7.4: A Tsetlin machine composite for breast cancer recurrence prediction.

Composite Decision. The normalised vote sums are then added together. In this way, you get the sum of votes for the composite as a whole. Here is the calculation for the breast cancer recurrence example:

$$\begin{aligned}\text{Composite Vote Sum} := {} & \text{Normalized Recurrence Vote Sum} + \\ & \text{Normalized Non-Recurrence Vote Sum.}\end{aligned} \tag{7.4}$$

The composite vote sum finally decides the output in the standard manner by comparing against zero:

$$\text{Predicted Class} \;:=\; \textit{Recurrence}\ \textbf{if}\ \text{Composite Vote Sum} \geq 0, \tag{7.5}$$

$$\textit{Non-Recurrence}\ \textbf{otherwise}. \tag{7.6}$$

Example of Composite Decision. Examine the input example and calculations in Figure 7.5. The Recurrence Tsetlin Machine's rules match this input, resulting in a vote sum of +2. The high vote sum indicates that the Tsetlin machine is confident. This vote sum is then normalized to +1 by dividing by the normalization factor 2. Oppositely, the Non-Recurrence Tsetlin Machine's rule does not match, giving the vote sum 0. This Tsetlin machine is uncertain. Consequently, the Recurrence Tsetlin Machine's normalized vote sum of +1 is the deciding factor, securing the Recurrence prediction.

Multi-Class Tsetlin Machine Composites. The Tsetlin machine composite in Figure 7.6 consists of four multi-class Tsetlin machines. The one on the left is based on a Histogram of Gradients, while the one next to it uses Colour Thresholding with 10x10 convolution filters. On the right side, the two Tsetlin machines use Colour Thermometers but with different convolution filters – 3x3 and 4x4.[2] Each of these Tsetlin machines has a different perspective on the input. The **argmax** operator decides the predicted class, which is the class with the highest vote sum.

[2]Convolution is covered in Chapter 4 while Chapter 9 addresses Image Understanding.

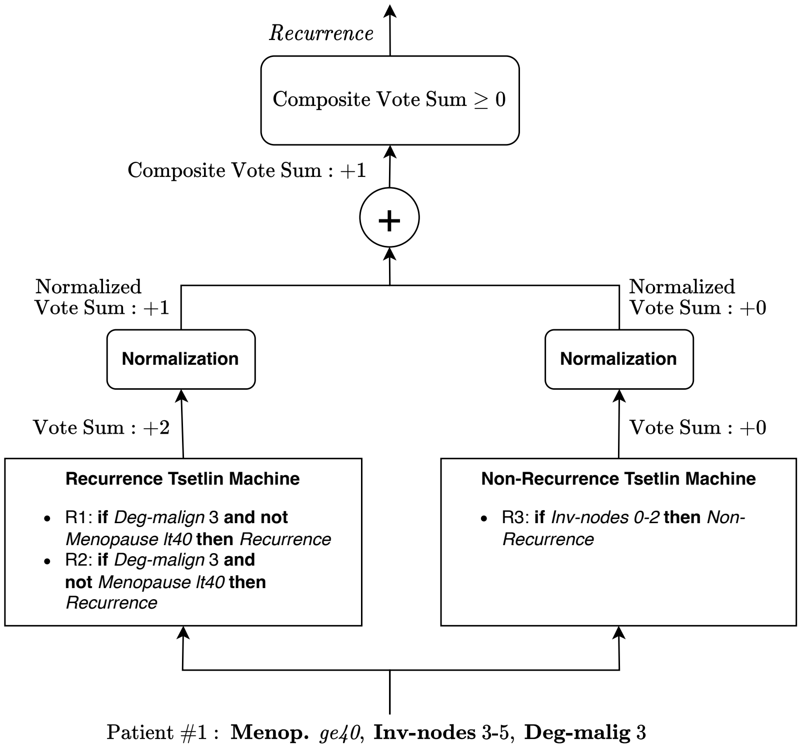

Figure 7.5: A Tsetlin machine composite for breast cancer recurrence prediction with example input and calculations.

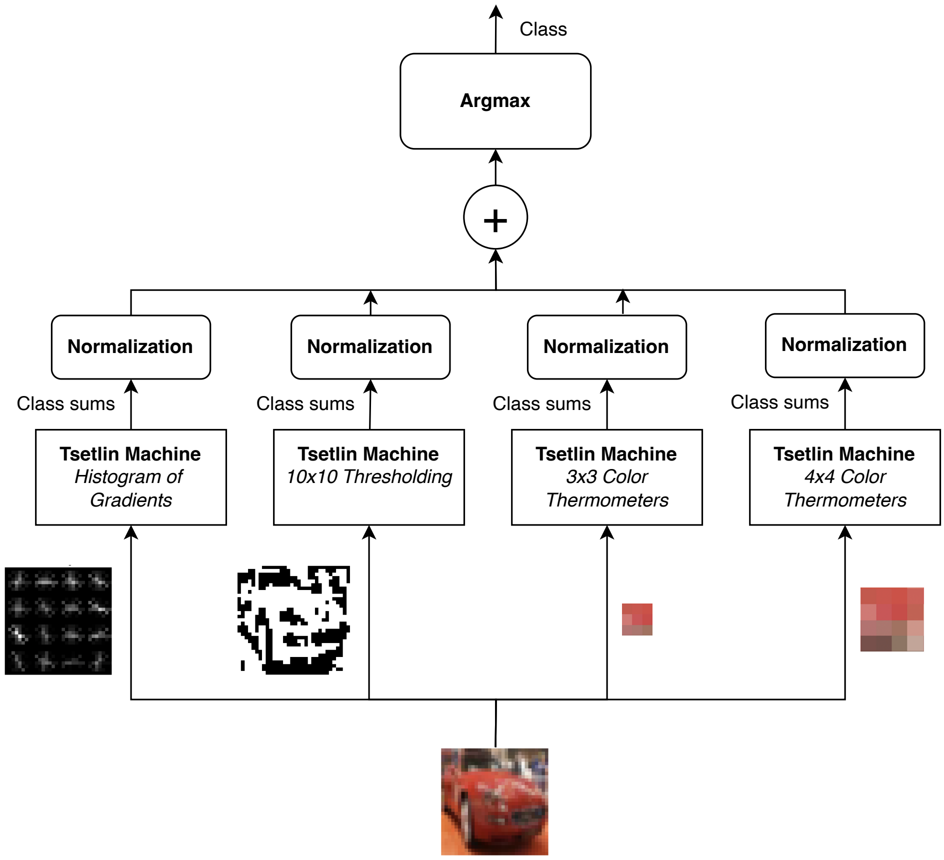

Figure 7.6: A multi-class Tsetlin machine composite for colour image classification.