Case Paper for District Level Deliberations on CT & AI 2026–27

Title of the Practice

From Play to Abstraction: Helping Children Understand Machine Learning, Neural Networks, Compression, and Large Language Models through Classroom Activities

Theme

Theme 2: From Play to Abstraction — Progressive Pedagogy for Grades 3–8

Target Group

This practice is designed for students of Grades 3–8. The same core idea can be used across this range with different levels of depth. For Grades 3–5, the focus is on observation, sorting, grouping, matching, guessing, and recognizing patterns through play. For Grades 6–8, the same activities are extended toward more abstract ideas such as features, training examples, classification, representation, compression, neural-network layers, prediction, and responsible use of AI.

The practice can be implemented in mathematics, science, computer science, language, social science, or interdisciplinary activity periods. It does not assume prior coding knowledge. The main requirement is curiosity, classroom participation, and the ability to observe patterns in familiar examples.

Core Context and Problem Statement

Artificial Intelligence is becoming a part of everyday life. Children see AI systems recognizing faces, recommending videos, answering questions, generating images, translating languages, and helping people search for information. However, most students encounter AI only as a finished tool. They see the output, but they do not understand the thinking process behind it. As a result, AI may appear magical, mysterious, or purely mechanical.

The central challenge addressed by this practice is: how can students in Grades 3–8 understand the basic ideas behind Machine Learning without beginning with advanced mathematics, coding, or technical definitions? The answer proposed here is to begin with the child's own experience. Every child already recognizes faces, animals, voices, letters, objects, and language patterns. Every child already groups things, makes guesses, learns from examples, and corrects mistakes. These natural learning processes provide a powerful entry point into Computational Thinking and AI.

This practice uses play-based and activity-based pedagogy to move students gradually from concrete experience to abstraction. It starts with a simple idea: the human brain is a pattern-finding system. Students first discover how they themselves recognize objects and make predictions. Then they compare this with how a machine learns from examples. After that, they explore how neural networks build larger patterns from smaller patterns. Finally, they connect the same principle to modern Large Language Models, which learn patterns in language and predict meaningful continuations from context.

The practice is not intended to make children experts in AI algorithms. Its purpose is to build conceptual readiness. Students should understand that AI is not magic. AI systems learn from examples, identify useful patterns, compress information into internal representations, and use those patterns to classify, predict, or generate outputs. At the same time, students should understand that AI systems can make mistakes when their examples are limited, biased, incomplete, or misleading.

CT and AI Competencies Addressed

This practice addresses several important Computational Thinking and AI-readiness competencies. Students learn pattern recognition by identifying repeated structures in images, objects, words, and sequences. They learn decomposition by breaking complex recognition tasks into smaller clues, such as shape, color, size, edges, parts, and context. They learn abstraction by ignoring unnecessary details and keeping only the useful features needed for recognition or prediction.

The practice also introduces classification, where students group new examples based on learned patterns. It introduces training and testing through the idea that a learner must first see examples and then apply the learned rule to new cases. It introduces generalization by showing that true learning is not just memorizing examples but using patterns on unseen examples. It introduces representation and compression by showing that both humans and machines cannot store every detail in a useful way; they must keep compact, meaningful information.

At a higher level, the practice introduces the idea of neural-network layers through physical classroom role play. Students understand that simple clues can be combined into larger patterns, and larger patterns can be combined into object recognition. The same foundation is then extended to language, where students explore how humans and language models predict likely words using context. The practice also includes responsible AI awareness by helping students understand that AI outputs are not always correct and must be checked thoughtfully.

Pedagogical Approach

The pedagogical approach follows the principle of moving from concrete play to abstract understanding. Instead of beginning with definitions such as “Machine Learning is a subset of Artificial Intelligence,” the teacher begins with questions that children can answer from experience. For example, students are asked how they recognize their friend’s face, how they know that a picture shows a dog, how they can read the same letter written in different handwriting, or how they can guess the next word in a familiar sentence.

The teacher then converts these familiar experiences into classroom activities. Students sort cards, complete patterns, train a classmate using labeled examples, create pixel art, act as layers of a human neural network, draw pictures from memory to understand compression, and predict missing words in sentences. Each activity is followed by discussion. The teacher gradually introduces vocabulary only after students have experienced the idea physically or visually.

This method is especially suitable for Grades 3–8 because young learners benefit from movement, play, visual material, group work, and storytelling. The abstraction is not forced at the beginning. It emerges naturally from the activity. Students first do the task, then describe what they did, and only then learn the formal idea behind it.

Classroom Implementation

The classroom implementation is organized as a sequence of six connected modules. These modules can be completed as a short workshop, a six-period classroom intervention, an AI-readiness club activity, or an interdisciplinary project across mathematics, science, and computer science.

Module 1: The Brain as a Pattern Finder

The first module begins with familiar visual and sensory examples. The teacher shows students partially hidden pictures, incomplete letters, blurred animals, simple shadows, or half-drawn faces. Students are asked to guess what they are seeing. In most cases, students are able to identify the object even when the information is incomplete.

The teacher then asks students how they were able to guess. Students usually respond that they had seen such things before, or that some clue helped them recognize the object. The teacher guides them to see that the brain compares the current input with past experience. It does not always need the full picture. It can use partial clues and still make a prediction.

This leads to the first major idea: a pattern is something useful that repeats across many examples. Students realize that they recognize a dog not because every dog looks exactly the same, but because many dogs share useful features. Similarly, they recognize the letter “A” even when different people write it differently because some underlying structure remains common.

At this stage, the teacher creates a classroom chart titled “My Brain Finds Patterns.” Students contribute examples such as recognizing a friend’s face, identifying a song from its tune, knowing a fruit from its smell, reading different styles of handwriting, or recognizing a vehicle from its shape. This chart becomes the foundation for the rest of the practice.

Module 2: Sorting, Grouping, and Classification

In the second module, students work with physical cards containing pictures of animals, fruits, vehicles, letters, shapes, leaves, and household objects. They are first asked to sort the cards into obvious groups. For example, they may group animals together, fruits together, and vehicles together. After this, the teacher asks them to group the same cards differently. They may sort by color, size, shape, number of legs, presence of wheels, living and non-living, round and non-round, or things found at home and things found outside.

This activity helps students understand that the same data can be organized in different ways depending on the feature selected. A mango can belong to the group “fruits,” but it can also belong to the group “yellow things,” “things with seeds,” or “things we eat.” Students begin to see that classification depends on the rule or feature being used.

The teacher then gives a new card and asks students where it should go. Students must justify their decision. This becomes the bridge to Machine Learning. A machine learning model also observes examples, identifies useful features, and then uses those features to classify new examples. For younger students, the word “feature” may be explained as “a useful clue.” For older students, the teacher may introduce the formal language of features, labels, classification, and prediction.

Module 3: Training a Human Model

The third module introduces the idea of training through role play. One student acts as a “machine.” The rest of the class trains this student using examples. For instance, the class may train the student to distinguish between “healthy food” and “junk food.” The class shows examples one by one: apple as healthy, chips as junk, carrot as healthy, cola as junk, milk as healthy, and burger as a more complex example depending on context.

The student acting as the machine must infer the rule from the examples. After the training round, the class tests the student with new examples such as banana, pizza, sprouts, fruit juice, chocolate, or homemade food. Sometimes the student-machine classifies correctly, and sometimes mistakes occur. These mistakes become important learning moments.

The teacher asks why the mistake happened. Students may say that not enough examples were given, the rule was not clear, some examples were confusing, or some items belonged to a grey area. Through discussion, students discover an important principle: learning from examples is powerful, but the quality and variety of examples matter.

This module introduces the difference between memorization and generalization. Memorization means remembering only the examples already seen. Generalization means applying the learned pattern to new examples. Students understand that a good learner, whether human or machine, should not merely memorize but should learn a useful rule or pattern.

Module 4: From Pixels to Pictures — How Machines See Images

The fourth module moves from human pattern recognition to machine image recognition. The teacher draws a square grid on the board and fills some squares to create a simple image such as a letter, a smiley face, a house, or an arrow. Students are then told that a computer does not see the picture exactly as humans do. It receives numbers. A filled square can be written as 1, and an empty square can be written as 0. The image becomes a grid of numbers.

Students create their own simple pixel-art images on graph paper. They then exchange their grids with other groups, who try to identify the image. This activity helps students understand that images can be represented as data. The teacher then shows two different versions of the same letter, such as two differently written “A”s. Although the exact filled squares may differ, students notice that some deeper structure remains similar.

This becomes the bridge to neural networks. The teacher explains that a neural network does not recognize a full object in one step. It builds understanding gradually. In simple terms, early layers may detect lines, edges, and curves. Later layers combine these into shapes. Still later layers combine shapes into parts. Finally, the model makes a prediction about the object.

For students, this can be enacted physically. Some students become the first layer and look only for simple clues such as lines, curves, colors, or corners. Other students become the second layer and combine these clues into parts. A final student gives the classification. This “human neural network” makes the layered nature of recognition memorable and concrete.

Module 5: Compression — Keeping What Matters

The fifth module introduces compression, one of the most important ideas in both human learning and Machine Learning. The teacher begins by asking students whether they remember every tiny detail of a face, an animal, or a classroom scene they saw yesterday. Students usually say no. The teacher then asks how they still recognize the person, animal, or place. Students respond that they remember important clues.

The teacher shows a detailed picture to the class for a short time and then removes it. Students are asked to draw what they remember. Most students do not reproduce every detail. Instead, they draw the main structure: the shape, key objects, important positions, or distinctive features. The class then compares what was preserved and what was forgotten.

This activity helps students understand that forgetting unnecessary details is not always a weakness. It can be useful. If the brain tried to remember every tiny detail of every image, sound, and experience, it would be overwhelmed. Instead, it keeps useful patterns. This is a child-friendly way to introduce compression.

The teacher explains that compression means keeping the useful information and reducing unnecessary detail. In AI, models also learn compact internal representations. A model trained on images does not simply store every image as a separate memory. It learns useful patterns that help it recognize new images. Similarly, a language model learns useful patterns in words, phrases, grammar, meaning, and context.

For older students, the teacher introduces the word “representation.” A representation is an internal form of information that keeps what is useful for a task. A face, a word, a sound, or an image can all be represented in a compressed form. This module helps students see that intelligence requires both pattern recognition and selective compression.

Module 6: From Image Recognition to Large Language Models

The final module extends the same ideas to language. The teacher writes incomplete sentences on the board, such as “The sun rises in the ***_,” “I opened the umbrella because it was*** _,” or “The teacher writes on the ____.” Students are asked to predict the missing word. Most students can guess likely words because they have heard and read many similar sentences.

The teacher then asks how they guessed the next word. Students realize that their brain uses context and past language experience. The teacher explains that Large Language Models also learn from huge amounts of text. They learn patterns in how words appear together, how sentences are structured, how meaning changes with context, and what continuation is likely.

The teacher carefully clarifies that a Large Language Model is not the same as the human brain. The comparison is used only to understand the broad idea of learning from examples and predicting from patterns. Humans have lived experience, intention, emotions, and understanding of the world in ways that machines do not. However, both humans and language models use previous patterns to make predictions.

The module closes by connecting all earlier activities. In image recognition, AI learns patterns in pixels. In speech recognition, AI learns patterns in sound. In language models, AI learns patterns in words and sentences. In all these cases, the broad process is similar: examples are observed, useful patterns are learned, information is compressed into internal representations, and the system uses those representations to classify, predict, or generate.

Sample Classroom Exercises

Several exercises are used throughout the practice. In the “Complete the Pattern” exercise, students continue simple sequences such as numbers, shapes, colors, or sounds. This builds pattern-recognition ability. In the “Guess the Object from Clues” exercise, the teacher gives clues one by one, and students identify which clue helped the most. This introduces feature importance.

In the “Train Your Friend” exercise, students give labeled examples to a classmate who acts as the machine. This helps them understand training, testing, and error correction. In the “Pixel Art Recognition” exercise, students convert drawings into grids of 0s and 1s, helping them understand that images can be represented as data. In the “Compression Drawing” exercise, students observe a detailed image and redraw it from memory, discovering that useful information is preserved while many details are ignored.

In the “Human Neural Network” exercise, students physically stand in layers. One group detects simple clues, another group combines clues, and the final group makes the prediction. In the “Predict the Next Word” exercise, students complete sentences and discuss how context helps prediction. Finally, in the “When AI Makes Mistakes” exercise, students explore ambiguous examples such as bat as an animal versus bat as sports equipment, or tomato as fruit versus vegetable. This helps them understand context, ambiguity, and responsible AI use.

Evidence-Based Impact

The impact of the practice can be measured through pre-activity and post-activity assessments, student worksheets, group charts, teacher observation, classroom discussion, and student reflection cards. Before the intervention, students may be asked simple questions such as: What is AI? How does a machine recognize a face? What is a pattern? What does it mean to learn from examples? Can AI make mistakes? What is the difference between remembering and understanding?

In many classrooms, students initially describe AI as a robot, a computer, or a smart machine. Very few students are likely to explain AI in terms of examples, patterns, classification, prediction, or errors. After the intervention, students should be able to use more meaningful explanations. For example, instead of saying “AI knows the answer,” students may say, “AI has learned from many examples and finds patterns.” Instead of saying “AI is always correct,” students may say, “AI can make mistakes if it learns from wrong or incomplete examples.”

The expected evidence of learning includes improved student ability to identify patterns, classify examples using features, explain training through examples, distinguish between memorization and generalization, describe an image as a grid of data, explain compression as keeping useful information, and connect word prediction with language models. Students should also show more responsible thinking by questioning AI outputs rather than accepting them blindly.

For final documentation, the school may collect before-and-after worksheets, photographs of group activities, samples of pixel-art grids, student-created classification charts, reflection cards, teacher observation notes, and short quiz results. A simple pre-test/post-test comparison may be used to show growth in conceptual understanding. For example, the number of students who can correctly explain “AI learns from examples” before and after the practice can be compared. Similarly, student explanations can be assessed qualitatively to show a shift from tool-based understanding to concept-based understanding.

Inclusivity

This practice is inclusive because it does not depend on expensive infrastructure, advanced devices, or prior coding knowledge. Most activities can be conducted with paper cards, charts, board drawings, physical movement, worksheets, and group discussion. Computer demonstrations may be added where available, but they are not essential for the core learning.

The design supports different learning styles. Visual learners engage through images, grids, charts, and drawings. Kinesthetic learners engage through sorting games, role play, and the human neural-network activity. Verbal learners engage through sentence prediction, explanation, and discussion. Analytical learners engage through classification rules and pattern discovery. Creative learners engage through drawing, storytelling, and designing their own examples.

The teacher can ensure inclusive participation by forming mixed groups, rotating leadership roles, inviting quieter students to explain observations, and ensuring gender-balanced participation in demonstrations. Since the activities begin with familiar everyday examples, students from different backgrounds can participate meaningfully. The focus is not on who already knows computers, but on how every child observes, thinks, classifies, predicts, and learns.

Scalability and Replication

The practice is highly scalable because it uses low-cost and easily available classroom materials. It can be replicated in urban, semi-urban, and rural schools. It can be implemented as a short AI-awareness module, a six-session classroom sequence, a mathematics-science integration activity, a computer science foundation unit, or a teacher-training demonstration.

The same framework can be adapted across subjects. In mathematics, it can connect with patterns, sequences, coordinates, shapes, symmetry, and graphs. In science, it can connect with classification of animals, plants, materials, body systems, and environmental data. In language, it can connect with word prediction, grammar patterns, sentence completion, and meaning from context. In social science, it can connect with map reading, grouping of resources, population patterns, and responsible technology use. In art, it can connect with visual patterns, symmetry, pixel art, and generative design.

The practice can also be scaled by preparing a simple teacher toolkit. The toolkit may include picture cards, sample grids, sentence prompts, classification worksheets, reflection questions, and a short teacher guide explaining how each activity connects to a CT or AI concept. Because the activities are modular, schools can adopt only some modules initially and expand later.

Challenges Faced

One possible challenge is that some teachers may initially feel that AI is too technical to teach in middle or preparatory grades. This challenge can be addressed by showing that the practice begins with everyday thinking rather than programming. Teachers do not need to explain complex algorithms. They need to guide students from familiar pattern recognition toward basic AI concepts.

Another challenge is avoiding oversimplification. The comparison between the brain and neural networks must be presented carefully. Students should understand that the brain and artificial neural networks are not identical. The analogy is useful for explaining pattern recognition, layered processing, and learning from examples, but it should not be presented as a complete scientific equivalence.

A third challenge is collecting strong evidence. Since the DLD format values demonstrated impact, teachers must plan evidence collection from the beginning. If the activity is conducted without pre-test, post-test, worksheets, reflection cards, or photographs, the practice may be strong pedagogically but weak in documentation. Therefore, evidence collection should be built into the activity design.

Learnings and Reflections

The most important learning from this practice is that AI can be introduced meaningfully to children when it is connected to their own thinking. Students do not need to begin with code or mathematical formulas. They can begin with recognizing faces, guessing objects, sorting cards, completing patterns, and predicting words. Once they experience these ideas physically and visually, abstract terms such as pattern, feature, classification, training, prediction, compression, and representation become easier to understand.

Another important learning is that compression is a powerful concept for children. When students draw a detailed picture from memory, they immediately understand that they do not store every detail. They preserve what matters. This gives them a simple but deep way to understand both human memory and machine learning representations.

The practice also shows that students learn better when they act out the system. When they become the data, the trainer, the model, the layer, and the evaluator, the idea of AI becomes less mysterious. The human neural-network activity is especially useful because it converts an abstract technical concept into a physical classroom experience.

A final learning is that responsible AI should be introduced from the beginning. Children should not learn only that AI is powerful. They should also learn that AI can be wrong, biased, incomplete, or confused by ambiguous context. This creates a balanced understanding of AI as a useful but imperfect technology that must be used thoughtfully.

Conclusion

This practice demonstrates a progressive pathway for introducing Machine Learning and AI readiness under the theme “From Play to Abstraction.” It begins with the child’s own brain as a pattern-recognition system and gradually moves toward sorting, classification, training from examples, image representation, neural-network layers, compression, and Large Language Models.

The central message for students is that AI is not magic. AI learns from examples, finds patterns, compresses useful information, and uses those patterns to classify, predict, or generate outputs. At the same time, AI can make mistakes, and human judgment remains important.

The practice is activity-based, inclusive, low-cost, interdisciplinary, and scalable. It helps students develop foundational CT and AI competencies without requiring advanced programming or expensive infrastructure. By connecting play, classroom charts, physical activities, simple visual representations, and reflective discussion, the practice prepares students to understand AI conceptually and responsibly.

Suggested Annexures

The following annexures may be attached to strengthen the final submission: pre-test and post-test worksheets, photographs of classroom activities, student-created classification charts, pixel-art grids, reflection cards, teacher observation notes, sample student explanations before and after the intervention, short quiz results, and feedback from students or teachers.

Sample Pre-Test Questions

What is a pattern? How do you recognize a dog even if every dog looks different? What does it mean to learn from examples? Can a machine make a mistake? What is AI? How does your brain guess the next word in a sentence?

Sample Post-Test Questions

How is classification connected to pattern recognition? What is the difference between memorization and generalization? Why does a machine need many examples to learn? How can an image be represented as numbers? What does compression mean? How is word prediction connected to Large Language Models? Why should we check AI answers before trusting them?

Sample Student Reflection Prompt

Today I learned that AI is not magic. It learns from examples and finds patterns. One activity that helped me understand this was _***. I also learned that AI can make mistakes when*** _. I will use AI responsibly by ____.