CODE THE CANVAS

AI-Powered Visual Systems for Creative Expression

Problem Statement

Every image communicates something before it is explained. Participants are challenged to design and develop an **AI-powered visual system** that brings this unspoken meaning to the surface. The system may use data and machine learning to generate, analyze, transform, or respond to images, producing outputs that convey a clear mood, story, or concept.

The objective is not merely to generate pictures. The system's **behaviour and technical design must express the idea** behind the work. Possible directions include artwork that changes with emotion, narratives that unfold scene by scene, or a visual canvas that responds interactively to its user.

Technical Expectations

Participants are expected to go beyond simply invoking a pre-trained model or external API. Each submission should demonstrate thoughtful system design, a transparent methodology, and a clear relationship between the technical approach and the intended creative outcome.

Teams should be prepared to explain the following aspects of their solution:

• **Data sources:** Identify the datasets, image sources, or other inputs used, and provide appropriate attribution.

• **Preprocessing:** Describe relevant cleaning, filtering, feature extraction, or data preparation steps.

• **Workflow:** Explain how inputs move through the system and how the final visual output is produced.

• **Prompting or fine-tuning:** Document the prompting strategy, model configuration, fine-tuning, or other adaptation methods, where applicable.

• **Original contribution:** Demonstrate an implementation, extension, or decision-making layer developed by the team.

• **Creative–technical connection:** Explain how the system's logic and behaviour support the mood, story, or concept being communicated.

Illustrative Example: "Mood Mirror"

Consider a system called **Mood Mirror**. A user uploads a photograph from their day—for example, a campus walkway at sunset. The system analyzes the image using OpenCV to extract visual properties such as dominant colours, brightness, and contrast. A pre-trained vision model such as CLIP may then estimate the image's mood, identifying it as "calm and nostalgic."

A diffusion model subsequently restyles the photograph as a painting using a warm, soft palette and loose brushstrokes. The team's original contribution is the documented mapping between extracted visual features and mood labels, and from those labels to prompt and style parameters. This mapping makes the

transformation explainable: the final image reflects the feeling the system inferred from the original.

Scope and Flexibility

The example above is provided for illustration only and is not a required solution. Teams are free to explore any theme, story, mood, or concept, provided that the system’s behaviour and output meaningfully reflect the chosen idea. Open-source tools and pre-trained models are welcome, provided they are properly credited and the team contributes its own implementation or extension.

Core principle: Build a visual system in which the technology does not simply create an image—it helps communicate an idea.