CENTRE FOR ENTERTAINMENT ARTS

WEEK 01

Storyboard I

AI Skills and Creative Practice

Lesson 1: Storyboard I

Directorial Pre-Production Intent · Directorial Prompt Construction · Panel Reading Against Intent

Lesson Objectives

1. By the end of this lesson, the learner will Position AI image generation as a directed pre-production move rather than an automated output
2. By the end of this lesson, the learner will Compose a single-panel prompt that carries cultural setting through visual detail
3. By the end of this lesson, the learner will Read a generated panel against the directorial intent that produced it

Lesson Learning Outcomes

- The learner can Distinguish directorial intent from a generated image
- The learner can Identify cultural register in a generated panel
- The learner can Evaluate a panel against directorial intent

INTRODUCTION — 10 MINUTES

Every AI image tool will generate a panel the first time you press the button. The panel will be composed, lit, and culturally plausible. It will also belong to the model, not to you — because the model filled every gap you left open. A generation that accepts the first output is not pre-production work. It is a handover.

Pre-production at any serious animation house works differently. Before a single frame is produced, a director specifies the visual outcome: what the panel must carry in composition, what cultural register must be legible without a caption, what the image has to do before it earns a place in the pitch. Nuqta Studio, a Riyadh-based animation house pitching short-form work for Saudi streaming platforms, cannot defend a panel to its writers' room unless someone made the directorial decision that produced it. The question is whether that someone was the maker or the model.

This course is about taking that decision back. Over the coming weeks you will build a directed storyboard pitch — panels generated through AI tools but controlled through prompt construction, iterated through structured logs, and compiled into a sequence that reads as authored work. The cultural setting you are working in, KSA-set scenes drawn from Saudi street life, family gatherings, and desert geography, cannot be left to model defaults. It has to be encoded into the composition, panel by panel, through every generation attempt.

That work starts here. The first question to answer is not how to use the tool — it is how to hold the directorial intent before you open it.

1 Directorial Pre-Production Intent

Position AI image generation as a directed pre-production move rather than an automated output

Mamoru Hosoda's production team on _Mirai_ (2018) documented a practice that set the film apart from contemporaries: before any scene entered production, the team made explicit decisions about framing, cultural object placement, and spatial register — decisions that were visible in the panels before any rendering pass began. The intent was legible in the composition. The panels were not discovered; they were directed.

Core: Directed Generation vs. Accepted Output

Pre-production AI image generation is a directed move when the maker holds a specific visual outcome in mind before the generator runs — and the generator is a tool executing that decision, not the source of decisions. When there is no prior intent, the generator fills every underdetermined gap with a model default: the most statistically probable cultural context, the most common compositional arrangement, the most frequent light-source placement. The output may be technically coherent and visually appealing, and it will carry none of the maker's direction.

Directorial intent is the specific visual outcome the maker is trying to achieve _before_ generation begins. It includes what the panel must show (composition, framing, spatial arrangement), what register it must carry (cultural specificity, tone, atmosphere), and — critically — what it must not require to communicate (a caption, a label, an explanatory title). A panel that needs a caption to locate its cultural setting has not been directed; it has been generated.

The gap between these two states is not about image quality. A first-output panel can be technically excellent. The gap is about authorship: who made the compositional decisions, the maker or the model?

Directed generation

Accepted output

Maker specifies visual outcome before generation

Maker evaluates output after generation

Cultural register is encoded in the prompt structure

Cultural register is left to model default

The panel is the result of a decision

The panel is the result of a probability

Rationale precedes the image

Rationale describes the image

Checkpoint

Look at two generated panels of the same scene — one produced by accepting a first output, one produced through explicit directorial choices. Which compositional elements in each panel can be traced back to a maker's decision, and which were filled by the model? How do you tell the difference?

REFLECTION ACTIVITY — 10 MINUTES

Write a pre-generation intent statement for your assigned KSA-set scene. The statement must name: what the panel must show in composition and framing, what cultural register it must carry without a caption, and at least one compositional element that the model will default away from if you do not specify it. The intent statement is the directorial test the panel you generate in the challenge will be evaluated against — not a description of what you hope to generate.

-

DISCUSSION — 10 MINUTES

In small groups, discuss "Directorial Pre-Production Intent". Share one example from your own experience or from work you have seen. Listen for differences in how peers interpreted the same material.

2 Directorial Prompt Construction

Compose a single-panel prompt that carries cultural setting through visual detail

The art-direction brief for _Cyberpunk: Edgerunners_ (Trigger/Netflix, 2022) constrained the generative visual language of Night City's Latinx district into a consistent spatial and chromatic register across an entire production — not by describing the aesthetic in general terms, but by specifying the exact compositional constraints that would force the right spatial and detail outcomes. The brief translated a cultural register into reusable syntax that worked at the level where the production could not afford to default.

Core: Prompt as Compositional Constraint

A directorial prompt encodes compositional constraints, spatial relationships, and cultural detail at the level where the model cannot default. It does not describe the scene — it specifies what the panel's structure must deliver.

The distinction is structural. A surface description gives the model material to work with and lets it decide how to arrange that material. A compositional constraint tells the model where elements sit relative to each other, what scale relationships hold, what spatial properties the frame must carry. The model still generates; the prompt makes the compositional space of acceptable outputs small enough that defaults cannot fill it.

Cultural visual detail is the class of detail that carries register through composition: architectural features that identify a specific cultural geography, clothing silhouettes that are culturally legible at a distance, light-source placement that carries a specific time and spatial quality, object specificity that names a thing rather than a category. "A market" is a description. "Stone arcade with alternating shadow bands at 2pm, foreground stall with stacked dried-goods sacks at waist height, single receding vanishing point" is a compositional constraint. The second prompt does not let the model decide what a market looks like.

```
Surface description (model fills gaps):
"a street market in Riyadh, warm light, busy"

Directorial prompt (model executes constraints):
"stone arcade, alternating shadow bands, midday overhead light,
foreground stalls at waist height, single vanishing point centre-frame,
crowd depth cued by overlapping silhouettes, no sky visible"
```

The test: read the output panel. Can someone locate the cultural setting in the composition without a caption? If not, the prompt is still a description.

Core: Cultural Visual Detail as Compositional Syntax

Translating a cultural register into prompt syntax requires moving from adjectives to spatial and compositional specifications. Adjectives describe properties; compositional syntax specifies arrangements. "Traditional" describes a property the model fills arbitrarily. "Pointed arch framing right edge, tiled geometric pattern at eye level, carved stone lintel above figure" specifies arrangements the model cannot default away from.

The move has three steps:

1. Identify the cultural visual detail that carries the register you need — the specific architectural feature, textile geometry, spatial quality, or named light condition.
2. Convert it from a described property to a spatial specification: where it sits in the frame, what scale relationship it holds, what the frame must show for it to be legible.
3. Remove the adjectives that give the model decision space. "Warm traditional Riyadh market" gives the model three degrees of freedom. "Riyadh suq arcade, stone vaulted ceiling, warm directional light from left arch opening, foreground brass vessels at knee height" removes them.

Checkpoint

Take the following surface description: "a Hejazi family gathering, evening, warm interior light." Rewrite it as a directorial prompt by converting at least two descriptive elements into compositional constraints. What specific cultural visual detail must the prompt name to carry the Hejazi register through composition rather than caption?

> **REFLECTION ACTIVITY — 10 MINUTES**

Using your Topic 1 intent statement as the constraint document, draft a directorial prompt for your KSA-set panel. The prompt must encode at least two compositional constraints derived from your intent statement, name at least one piece of cultural visual detail specific enough that the model cannot default to a generic substitute, and pass the caption-free test: a reader of the generated panel should be able to locate the cultural setting in the composition without reading any accompanying text.

-

DISCUSSION — 10 MINUTES

In small groups, discuss "Directorial Prompt Construction". Share one example from your own experience or from work you have seen. Listen for differences in how peers interpreted the same material.

3 Panel Reading Against Intent

Read a generated panel against the directorial intent that produced it

Jonathan Glazer's _The Zone of Interest_ (2023) is built on a documented directorial intent: to render the domestic normalcy of a household adjacent to atrocity without recourse to dialogue or caption as the primary carrier of meaning. Glazer's compositional decisions — framing that keeps the horror outside the frame's edge, spatial arrangement that treats the ordinary as the subject — are analysable against that intent at the level of framing and spatial arrangement alone. The film's critical practice demonstrates that compositional choices carry directorial meaning in a way that is readable by a viewer who holds the intent and looks at the composition.

Core: Panel Reading as Directorial Evaluation

Evaluating a panel against directorial intent means asking whether the compositional choices in the panel — framing, spatial detail, register — deliver the specific outcome the direction specified. It is not aesthetic judgment. It is not a question of whether the image is good. It is a question of whether the panel passes the directorial test the intent statement set.

The reading practice has a specific sequence:

1. State the directorial intent: what was the panel supposed to show, carry, and communicate without a caption?
2. Identify which compositional elements in the panel carry that intent — which spatial arrangements, which detail choices, which framing decisions are legible as directed moves.
3. Identify which elements are model defaults that the direction did not overcome — the generic compositional choices the model made in the absence of a sufficient constraint.
4. Name the gap: what specific revision to the prompt would address the default?

A panel that requires caption text to locate its cultural setting has failed the directorial test regardless of image quality. The failure is not an aesthetic failure; it is an authorship failure. The model directed that element, not the maker.

Core: Locating Directorial Failure in the Composition

Model defaults are consistent and identifiable. A model asked to generate a "market scene" without spatial constraints will produce a wide-angle establishing shot with a crowd receding to a vanishing point and colourful foreground stalls — because that is the most statistically probable compositional arrangement for the category "market." When that arrangement appears in a panel whose directorial intent specified a specific spatial register, the default is evidence that the prompt did not constrain that decision.

Reading for defaults means naming what the model decided rather than the maker. This is not a post-hoc rationalisation of what the panel looks like; it is a comparison between the intent statement and the composition. The gap between them is the revision agenda for the next generation attempt — the same cognitive move the iteration record in the following week will require.

Checkpoint

Present a generated panel to yourself with no caption and no visible prompt. State: what cultural setting does this panel carry? Where in the composition — not in any accompanying text — is that setting legible? Which compositional elements are traceable to a directorial decision and which are model defaults? What would you revise in the prompt to address the defaults?

REFLECTION ACTIVITY — 10 MINUTES

Apply the reading practice to the panel you generated in Topic 2. Write a one-paragraph evaluation (four to six sentences) that states: what the directorial intent specified, what compositional elements in the panel deliver that intent, which elements are model defaults the direction did not overcome, and whether the panel passes the caption-free test. The evaluation must name specific compositional elements — not aesthetic impressions — as evidence for each claim.

DISCUSSION — 10 MINUTES

In small groups, discuss "Panel Reading Against Intent". Share one example from your own experience or from work you have seen. Listen for differences in how peers interpreted the same material.

HOMEWORK

Part 1: Observe the Full Directorial Move

Your instructor will demonstrate the complete decision chain for a non-KSA scene: stating directorial intent aloud, writing a prompt that encodes compositional and cultural-detail constraints rather than surface descriptions, generating a panel, and reading the output against the stated intent. Watch for the specific moves — where constraints replace adjectives, which cultural visual detail gets named, and what the caption-free test reveals.

Part 2: Write a Pre-Generation Intent Statement

Choose a KSA-set scene for your panel. It can be a Riyadh street scene, a Hejazi family interior, a desert journey, or another Saudi-set scenario you know visually.

Write a pre-generation intent statement covering three things:

1. What the panel must show in composition and framing — name the spatial arrangement, not the subject matter.
2. What cultural register the panel must carry — name at least one piece of cultural visual detail that must be legible in the composition without caption text.
3. At least one compositional element the model will default away from if you do not specify it — name the default and name your constraint.

Bring this statement to a peer or instructor. Their test is simple: is the statement specific enough to generate against, or does it still leave the model room to decide? Revise until it passes that pressure.

Part 3: Write a Directorial Prompt and Generate

Convert your intent statement into a directorial prompt. Replace every adjective that gives the model decision space with a spatial or compositional specification. Name cultural visual detail at the level where the model cannot default to a generic substitute.

Generate your panel using Midjourney, Stable Diffusion via a hosted interface, or another AI image tool.

Apply the caption-free test: can a reader locate the cultural setting in the composition without reading your prompt or any caption text? If not, identify which constraint the prompt failed to encode, revise the prompt, and regenerate. You may make as many generation attempts as needed — document only the final prompt alongside the panel you submit.

Part 4: Write a Rationale

Write a short rationale (three to five sentences) defending the panel against your intent statement. The rationale must:

- Name the compositional elements in the panel that carry the directorial intent.
- Identify at least one element that is a model default the direction did not fully overcome.
- State whether the panel passes the caption-free test and why.

The rationale is not a description of what the panel looks like. It is a comparison between what the intent specified and what the composition delivers.

Assemble and Submit

Collect the following four items into a single folder and compress as a zip file:

- A plain-text file containing your pre-generation intent statement, named `intent-statement`
- A plain-text file containing the directorial prompt you submitted for the final generation, named `prompt`
- Your generated panel at its native resolution (`panel.png` or `panel.jpg`)
- A plain-text file containing your three-to-five sentence rationale, named `rationale`

Submission
Submit the zip file containing your intent statement, directorial prompt, generated panel, and rationale.

Feedback questions

- Which element of your intent statement was hardest to encode as a compositional constraint, and what made that translation difficult?
- Where in the panel can you locate a model default that your prompt did not overcome, and what specific constraint would you add to address it in the next attempt?

References

Readings

Prompting AI Art — The Creative Skill of Prompt Engineering — Oppenlaender et al. establish that prompt engineering is a non-intuitive, acquirable skill requiring style-specific vocabulary — directly framing why directorial prompt construction is learnable rather than intuitive.
https://arxiv.org/abs/2303.13534

Watch

Practitioner Talk — AI in Pre-Production [CURATOR NEEDED] — A short clip from a storyboard artist or production director describing their visual intent before a generation pass — not a tool tutorial; a SIGGRAPH or FMX practitioner talk on AI in pre-production would serve this slot and requires human curation before publication.
https://www.siggraph.org/learn/conference-collection/

Play

Midjourney Prompt Basics — Read the prompt syntax rules then write and generate a test panel, entering the tool's decision surface before the lesson models the directorial move.
https://docs.midjourney.com/hc/en-us/articles/32023408776205-Prompt-Basics

Resources

Midjourney Parameter List — Reference for aspect ratio, style weight, character reference, and compositional constraint parameters — the specific levers that encode directorial decisions into a generation call.
https://docs.midjourney.com/hc/en-us/articles/32859204029709-Parameter-List

Low-Stakes Retrieval Practice

Recall Test

Based on the content of this lesson

Your instructor will conduct a retrieval practice recall test based on the content of this lesson.

The test is low-stakes · it is not formally graded. Its purpose is to strengthen your long-term memory of the material by requiring you to actively retrieve what you have learned, rather than simply re-reading it. Research shows this is one of the most effective study strategies available.

What to expect:

- 15 minutes, closed notes
- Questions will draw directly from this lesson's concepts, vocabulary, and processes
- You will receive immediate feedback · the goal is to identify gaps, not to catch you out

Why retrieval practice works

The Testing Effect

Roediger & Karpicke (2006) demonstrated that students who practise recalling material retain significantly more over time than those who spend the same time re-studying.

Desirable Difficulties

Bjork (1994) showed that learning strategies that feel harder in the short term · such as recall · produce stronger, more durable long-term memory.

Principles of Instruction

Rosenshine (2012) identifies daily review and retrieval as one of the most evidence-backed habits of effective instruction.

Roediger, H. L. & Karpicke, J. D. (2006). Psychological Science, 17(3), 249-255.
Bjork, R. A. (1994). In Metcalfe & Shimamura (Eds.), Metacognition. MIT Press.
Rosenshine, B. (2012). Principles of Instruction. American Educator, 36(1), 12-19.