



China's Revised Patent Examination Guidelines for AI: Six Practice Tips for Drafters

By [*Charles W. Gray*](#)

Effective January 1, 2026, CNIPA's revised Patent Examination Guidelines add AI-specific examination rules in three layers: an ethics and legal compliance threshold that may be considered independently of novelty and inventive step; a higher inventive step bar that rejects mere scenario migration; and a stricter sufficiency of disclosure standard that materially increases the risk of black-box AI drafting. For US drafters filing into China, the mechanical changes are small but the drafting posture shifts materially. Six practice tips to conform now.

Tip 1. Add a PIPL and fairness compliance statement to every AI-adjacent spec.

The new Section 6.1.1 of Chapter IX, Part II makes ethical and legal compliance a threshold patentability issue that may be considered independently of novelty and inventive step. Covert facial recognition and demographic-weighted decision logic have been named in the illustrative cases as examples likely to be rejected where the claimed invention violates applicable legal, ethical, or public-interest requirements. When drafting for China, add a short compliance recital covering (a) lawful data acquisition under the Personal Information Protection Law with separate user consent where required, (b) anonymization or de-identification of personal data, and (c) controls designed to avoid unlawful discriminatory treatment or other legally prohibited outcomes in any decision-making, recommendation, or scoring model. Keep the recital in the specification so it travels with any US continuation applications and any Chinese divisional applications.

Tip 2. Stop leading with the application scenario.

The new inventive step examples confirm that transplanting a known deep learning model to a new field, without more, no longer supports inventive step in China. A patent application that pitches "conventional CNN, new target object" is likely to face rejection. Frame the disclosure around targeted modifications to model architecture, layer

configuration, loss function, training pipeline, or data-processing workflow, and describe the technical problem those modifications solve. The application scenario supports the story; it may no longer be the story.

Tip 3. Disclose the model at module resolution, not block-diagram resolution.

Section 6.3.1 now provides that the specification should record, where necessary to enable the invention and support the asserted technical contribution, essential modules, hierarchical structures, connection relationships, training procedures, and necessary parameters. A one-page architecture figure with three labeled boxes is not enough. For each learned component, disclose what it does, how it is trained, what data feeds it, and how it interconnects with adjacent modules. Well-known conventional building blocks may be described briefly, but any element that carries inventive weight has to be laid out in enough detail for a person skilled in the art to reproduce.

Tip 4. Explain the input to output causal logic.

The new “black box” example rejects specifications that state, in substance, “feed features A, B, and C to the model to obtain output D” without connecting the inputs to the target technical problem. When drafting, disclose (a) why each input feature is selected, that is, the causal link between the feature and the technical result; (b) what each network module does with the feature; and (c) how the model output maps to the ultimate technical effect. This is primarily a specification and sufficiency issue, though it may also affect support for the claimed scope; addressing it up front avoids sufficiency of disclosure attacks that are hard to cure by amendment.

Tip 5. Use permissive framing to protect claim breadth on non-inventive details.

Where a detail is not inventive but has to be disclosed to satisfy sufficiency, use permissive phrasing that a person skilled in the art may resolve as a routine design choice, such as “in one embodiment,” “in some embodiments,” or “may be configured as.” However, exercise caution with explicit statements that a given module’s position or hyperparameter falls within routine design choices in the field; while such statements can help keep those details out of the claim scope without inviting a sufficiency of disclosure objection, direct admissions that features are “routine” could be cited against inventive step if those features later become important. Reserve “routine design choice” language for details that are clearly not part of the inventive contribution. This mirrors the disclosure conventions we already use in US practice.

Tip 6. Quantify the technical effect.

The revised inventive step framework rewards applications that show a measurable delta over prior art. Include comparative experimental data, ablations, and quantitative metrics (accuracy, precision or recall, computational overhead, memory footprint, robustness under distribution shift, latency) that demonstrate the effect of the disclosed architectural or training changes. Numeric evidence can support the inventive step argument where prior art has no comparable data, and it can help defend the record against a routine-optimization rejection.