

Generative AI Workflows for Visual Effects:

Integrating ComfyUI™ with NUKE®

by Victor Perez

www.victorperez.co.uk

FOUNDRY Certified Trainer.

Victor Perez

Comfy

COURSE DESCRIPTION

In an era where visual storytelling is evolving at light speed, **Generative AI Workflows for Visual Effects: Integrating ComfyUI with NUKE** offers a new path forward for artists ready to embrace the creative revolution. Tailored for visual effects professionals, compositors, and digital creatives, this course is designed to bridge the gap between cutting-edge AI tools and the traditional artistry of film and visual narrative. Whether you're an industry veteran or an emerging artist, this course provides the essential knowledge and practical skills to future-proof your workflow, and your career.

This is not just a software tutorial. It's a holistic, structured journey through six powerful domains that define the present and future of creative production. We begin with **Artificial Intelligence Fundamentals**, where you'll gain a solid understanding of concepts like *machine learning*, *neural networks*, and *stable diffusion*, presented in a software-agnostic way to build a foundation that transcends platforms. You don't need to come in with any prior experience in AI; this course welcomes you as you are and builds you up from the ground level.

Next, we tackle **Hardware Explained**, a vital section that demystifies the technical requirements of working with AI locally or in the cloud. From understanding GPUs to optimising your setup for performance and affordability, this part of the course empowers you to make smart decisions that support your creative vision.

In a world where technology moves faster than legislation, we must stay grounded in our values. That's why we dedicate time to **Ethics and Responsibilities**, because generative AI is a powerful tool that must be wielded with awareness, respect, and integrity. Together, we'll explore the implications of authorship, copyright, bias, misinformation, and responsible sourcing, helping you become not just a better artist, but a conscious one.

At the heart of the course lies **Using ComfyUI**, a deep and detailed exploration of one of the most powerful and flexible open-source interfaces for generative AI. You'll learn how every node works, how to build workflows, understand metadata, optimise processing, and use tools like ControlNets, LoRAs, and upscaling techniques. From prompt-based image generation to image-to-image transformations and even AI-driven video workflows, you'll master every corner of the tool, and understand how it all connects back to the core principles of diffusion models.

Then we dive into **NUKE + ComfyUI**, where the real magic happens. If you're already familiar with NUKE, this section will teach you how to supercharge your compositing workflow with generative AI, using ComfyUI not just as a creative sandbox but as a professional-grade VFX engine. You'll learn roundtrip workflows, image integration, and how to apply AI as a tool of precision, not a replacement. You'll become the kind of artist studios need: fast, imaginative, and technically fluent.

Finally, we open the door to the **Community Tools & Resources** that are fuelling this open-source explosion. You'll learn where to find trusted models, how to contribute back, and how to stay updated in a fast-moving ecosystem. We'll give you documentation, references, and platforms to expand your knowledge well after the course is done.

But above all, this course is designed to empower you. AI is not the end of the artist, it's a new beginning. If you feel stuck in your career, this may not be your fault. The industry is shifting. This is the biggest leap since the transition from optical to digital in the 1990s. And like then, those who adapt will lead.

You'll learn through a series of pre-recorded tutorials by one of the most influential VFX trainers in the world, Victor Perez, who has been training a whole generation of VFX artists for the las 15 years. The course will be completed with optional live sessions, ensuring you can study at your own pace while staying connected to a community of like-minded creatives. The course is offered in English, Spanish, Italian, Hindi, and Chinese, making it truly global and inclusive.

This is your moment to stop fearing AI and start using it. This is the course that gives you everything you need to master ComfyUI, understand the technology behind it, and apply it to the art of cinematic storytelling. From beginner to advanced, all in one place.

Welcome to the next generation of VFX. Welcome to the future of your craft.

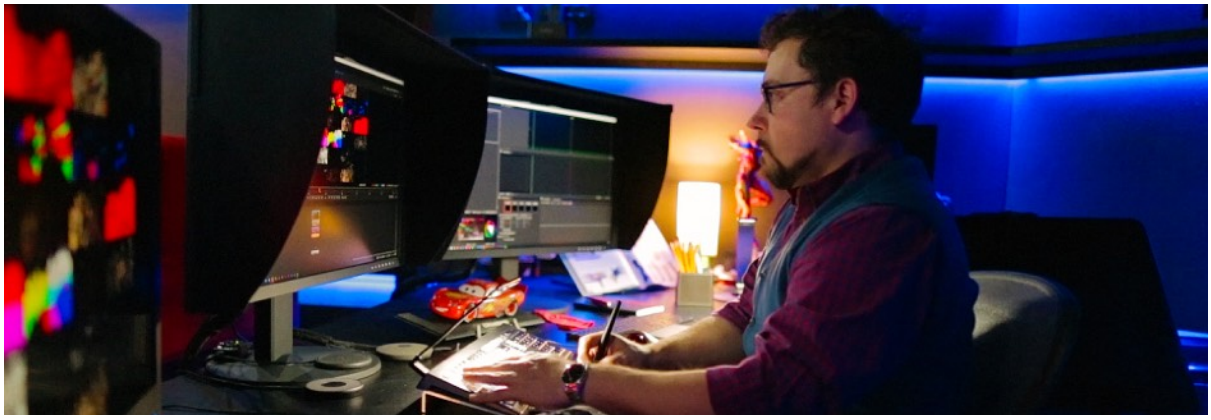

COURSE SYLLABUS

1. Artificial Intelligence Fundamentals
1.1. What is AI?
1.2. Foundations of Machine Learning
1.2.1. Machine Learning vs Deep Learning
1.2.2. Inference
1.2.3. Epoch
1.2.4. Batch Size
1.2.5. Steps
1.2.6. Weights
1.2.7. Overfitting vs Underfitting
1.3. Neural Networks & Architectures
1.3.1. NN (Neural Network)
1.3.2. CNN (Convolutional Neural Network)
1.3.3. Transformers
1.3.4. LLM (Large Language Model)
1.3.5. GPT (Generative Pre-trained Transformer)
1.3.6. Visual Transformers
1.4. Core Technologies & Optimization
1.4.1. Tensor
1.4.2. Data Types & Precision
1.4.3. CUDA & GPU Acceleration
1.4.4. TensorRT & ONNX
1.4.5. Introduction to LoRAs (Low-Rank Adaptation)
1.4.6. Quantisation
1.4.7. Training vs Inference Optimisation
1.4.8. GGML (Georgi Gerganov's Machine Learning)
1.4.8.1. GGUF (Grokking GGML Unified Format)
1.5. Diffusion Fundamentals
1.5.1. History of Diffusion Models: Mathematics and Metaphors
1.5.2. Diffusion
1.5.2.1. SD (Stable Diffusion) Architecture
1.5.2.2. UNet Architectures
1.5.2.3. DiT (Diffusion Transformer)
1.5.2.4. SD Checkpoint
1.5.2.5. Latent Image
1.5.2.6. Pseudo-Randomisation
1.5.2.6.1. Seed
1.5.2.7. VAE (Variational Autoencoder)
1.5.2.7.1. The Role of VAE in Diffusion Models
1.5.2.8. Latent Space Manipulation
1.5.2.9. Denoise
1.5.2.11. Conditioning
1.5.2.11.1. Prompts
1.5.2.11.2. Image Conditioning
1.5.2.11.3. Positive and Negative Conditioning
1.5.2.12. ControlNets
1.5.2.12.1. Pose
1.5.2.12.2. Depth Maps
1.5.2.12.3. Edge Maps
1.5.2.12.4. Normal Maps
1.5.2.12.5. Scribbles / Sketches
1.5.2.12.6. Segmentation Maps
1.5.2.12.7. OpenPose Skeletons
1.5.2.12.8. Line Art, QR Codes, Logos, etc.
1.5.2.13. LoRA (Low-Rank Adaptation)
1.5.2.14. TeaCache (Timestep Embedding Aware Cache)
1.6. Diffusion Sampler and Scheduling
1.6.1. Diffusion Sampler
1.6.2. Diffusion Scheduler (a.k.a. Beta Scheduler)
1.6.3. CFG (Classifier-Free Guidance)
1.6.4. Sampling Algorithms
1.6.4.1. Euler Family
1.6.4.2. DPM Family (Denoising Diffusion Probabilistic Models)
1.6.4.3. DPM++ Family (Improved Numerical Methods)
1.6.4.4. UniPC Family (Unified Predictor-Corrector)
1.6.4.5. LCM and Derivatives (Latent Consistency Models)
1.6.4.6. Classic Legacy Samplers
1.6.4.7. Modern Solver-Based Samplers
1.6.4.8. Experimental, Academic & Research Samplers
1.6.4.9. Special Purpose / Seed Variation Tools
1.6.4.10. Conclusions about Sampling Algorithms
1.6.5. Diffusion Time Step & Noise Schedule Visualisation
1.6.6. Beta Schedulers
1.6.6.1. Simple / Linear-Based Schedulers
1.6.6.2. Mathematical Curve-Based / Parametric Schedulers
1.6.6.3. Theory-Driven / Research-Based Schedulers
1.7. Model Architectures & Categories
1.7.1. Stable Diffusion Models
1.7.2. Model Types
1.7.2.1. Creative Generation (Idea to Image)
1.7.2.1.1. Text-to-Image
1.7.2.1.2. LoRA-Based Style Injection
1.7.2.1.3. Prompt-Based Variation Generation
1.7.2.1.4. Chaining Multiple Models for Refinement
1.7.2.2. Structure-Guided Generation (Form to Image)
1.7.2.2.1. Image-to-Image (img2img)
1.7.2.2.2. Depth-to-Image
1.7.2.2.3. Sketch-to-Image
1.7.2.2.4. ControlNet-Guided Generation
1.7.2.3. Restoration and Repair
1.7.2.3.1. Inpainting
1.7.2.3.2. Outpainting
1.7.2.3.3. Latent Space Manipulation
1.8. Upscaling & Super-Resolution
1.8.1. Traditional Interpolation vs AI Generation
1.8.2. Types
1.8.2.1. ESRGAN (Enhanced Super-Resolution GAN)
1.8.2.2. Real-ESRGAN
1.8.2.3. SwinIR (Shifted Window Transformer for Image Restoration)
1.8.2.4. Custom-Tuned Upscaler Models
1.8.2.4.1. UltraMix Balanced
1.8.2.4.2. Remacri
1.8.2.4.3. 4x-UltraSharp
1.8.2.4.4. CleanCrop / RestoreFormer
1.8.2.5. Latent Upscaling
1.8.2.6. SeedVR2
1.8.3. Popular Uses
1.9. Generative Model Families for Visual Synthesis
1.9.1. Autoregressive Models
1.9.2. GANs (Generative Adversarial Networks)
1.9.3. VAEs (Variational Autoencoders)
1.9.4. Diffusion Models
1.9.4.1. Latent Diffusion Model (LDM)
1.9.5. Transformers for Vision
1.9.6. Hybrids
1.9.7. Generative Model Families
1.9.7.1. Foundational Stable Diffusion Lineage
1.9.7.1.1. Stability AI
1.9.7.1.2. Stable Diffusion 1.x
1.9.7.1.3. Stable Diffusion 2.x
1.9.7.1.4. Stable Diffusion XL 1.0 (SDXL 1.0)
1.9.7.1.5. SDXL Turbo & Lightning
1.9.7.1.6. Stable Cascade
1.9.7.1.7. Stable Diffusion 3.0 / 3.5
1.9.7.2. Community Developed Models
1.9.7.3. Community Developed Models
1.9.7.3.1. Realistic Vision
1.9.7.3.2. Real Dream Pony
1.9.7.3.3. Analog Diffusion
1.9.7.3.4. Analog Madness

1.9.7.3.6. OpenJourney
1.9.7.3.7. DreamShaper
1.9.7.3.8. Deliberate
1.9.7.3.9. Union Pro (Unified ControlNet Checkpoint)
1.9.7.4. Open-Source Multimodal Transformers
1.9.7.4.1. DeepSeek Janus Pro
1.9.7.4.2. SAM (Segment Anything Model) by Meta
1.9.7.4.3. VACE (All-in-One Video Creation and Editing)
1.9.7.4.4. LTX-Video (Latent Temporal eXpansion)
1.9.7.4.5. FLUX by Black Forest Labs
1.9.7.4.6. Master Flux Kontext (Context-Aware Image Editing)
1.9.7.4.7. Wan (Image-to-Video Workflow)
1.9.7.4.8. HiDream
1.9.7.4.9. OmniGen (Unified Image Generation)
1.9.7.4.10. Qwen-Image
1.9.7.5. Other Independent, Commercial, and Research Models
1.9.7.5.1. HunyuanVideo by Tencent
1.10. Model Configuration and Optimization
1.10.1. Model Selection and Classification
1.10.1.1. Model Configuration Selectors
1.10.1.2. Speed / Performance-Based
1.10.1.3. Experimental / Development
1.10.1.4. Style / Training Variant
1.10.1.5. Compatibility Tags
1.10.1.6. Internal Precision or Size
1.10.2. API (Application Programming Interface) Oriented Models
1.10.3. Optimization Principles and Best Practices
1.10.3.1. Interpret Performance Metrics
1.10.3.2. Accuracy vs Speed
1.10.3.3. Choose Optimal Tensor Formats
1.10.3.4. Use GPU Resources Wisely
1.10.3.5. Accelerate Inference Using Quantised Models
1.10.3.6. Training vs Inference Optimization
1.10.3.7. Recommended Settings and Use Guidelines
1.11. Prompt Engineering for Generative AI
1.11.1. What Is a Prompt?
1.11.2. What Is Prompt Engineering?
1.11.3. Key Elements of an Effective Prompt
1.11.4. Best Practices for Prompt Creation
1.11.5. Common Prompt Engineering Tools
1.12. Prompt Engineering Techniques and Advanced Approaches
1.12.1. Text-to-Text Prompt Techniques
1.12.2. The Interview Pattern Approach
1.12.3. The Chain-of-Thought Approach
1.12.4. The Tree-of-Thought Approach
1.12.5. Multimodal Prompting
1.12.6. The Playoff Method
1.12.7. Choosing the Right Prompting Approach
1.13. Conclusion

2. Hardware Explained

2.1. CPU (Central Processing Unit)
2.1.1. The Role of the CPU in AI workflows
2.1.2. Multithreading
2.1.3. Instruction Sets (AVX, AVX-512)
2.1.4. When CPU Speed Matters vs When It Is Secondary to GPU Performance
2.1.5. ECC Support for High-Reliability Workflows
2.1.6. Emerging Trends
2.1.6.1. Apple Silicon M-Series with Integrated Neural Engines
2.1.6.2. AMD's AI Extensions
2.1.6.3. Intel's Gaudi Accelerators and AI Push
2.1.6.4. Hybrid Architectures and On-Die AI Cores
2.2. GPU (Graphics Processing Unit)
2.2.1. The Role of the GPU in AI workflows
2.2.2. CUDA Cores
2.2.3. Tensor Cores
2.2.4. CUDA Cores vs Tensor Cores
2.2.4.1. ROCm Ecosystem Maturity for AMD GPUs as an Alternative to CUDA
2.2.5. VRAM Architecture
2.2.6. GPU Memory Bandwidth
2.2.7. Clock Speeds
2.2.8. Multi-GPU Setups
2.2.8.1. NVLink
2.2.8.2. PCIe Lane Sharing
2.2.8.3. Scaling Limitations
2.3. VRAM (GPU Memory)
2.3.1. Why VRAM Size Is Critical for Large Models and High-Resolution Image Generation
2.3.2. VRAM Speed
2.3.3. VRAM Memory Types
2.4. RAM (System Memory)
2.4.1. Recommended Sizes for AI Workflows
2.4.2. ECC RAM for Data Integrity
2.4.3. How RAM Speed Interacts with CPU and Storage
2.5. Storage
2.5.1. NVMe vs SATA SSD vs HDD
2.5.2. Read/Write Speeds and Caching Behaviour
2.5.3. Data Transfer Pipelines: From Storage to GPU
2.6. AI Accelerator Cards / Dedicated Chips
2.6.1. What They Are and When They Are Worth It
2.6.2. Limitations in Artist Workflows
2.6.3. Specialised Inference Hardware Like NVIDIA Jetson for Edge AI
2.7. Cloud vs Local Hardware – Introduction
2.7.1. Latency
2.7.2. Cost Over Time
2.7.3. Scalability
2.7.4. Networking Requirements for Hybrid Setups
2.7.5. Data Security & Control Over Assets
2.8. Compatibility & the CUDA Ecosystem
2.8.1. Driver Versions
2.8.2. Framework Compatibility
2.8.3. Cross-Platform Considerations
2.9. Thermal Management & Power Requirements
2.9.1. Cooling Solutions: Air vs Liquid Cooling
2.9.2. VRM Quality
2.9.3. Fan Curves
2.9.4. PSU Wattage and Stability for Sustained AI Loads
2.9.5. Avoiding Thermal Throttling and Voltage Drop
2.9.6. Form Factor & Expansion Slots
2.10. Interconnect Standards
2.10.1. PCIe Gen4 vs Gen5 Performance Implications
2.10.2. Future CXL (Compute Express Link) for Shared Memory Between Devices
2.11. Energy Efficiency vs Performance
2.11.1. Impact of Power Draw on Cost and Sustainability
2.11.2. Efficiency Tuning for Mobile and Desktop Systems
2.12. Benchmarking & Bottleneck Identification
2.12.1. GPU/CPU/RAM/Storage Benchmarking Tools
2.12.1.1. Synthetic Benchmarks
2.12.1.1.1. Geekbench
2.12.1.1.2. CUDA-Z
2.12.1.2. AI-Specific Benchmarks
2.12.1.2.1. MLPerf
2.12.1.2.2. ComfyUI Test Graphs
2.12.1.2.3. Downloading ComfyUI Test Graphs
2.12.2. Recognising Real-World Bottlenecks
2.12.2.1. RAM / VRAM Overflows
2.12.2.2. Symptoms of Performance Issues in Multi-Node Pipelines
2.13. Turnkey Solutions and Acquisition Guidelines
2.13.1. How to Choose the Right Machine for Your Workload
2.13.2. Balancing Cost, Scalability, and Compatibility
2.13.3. Power Backup & Reliability for Uninterrupted Training/Rendering
2.13.3.1. UPS Systems (Uninterruptible Power Supplies)
2.13.3.2. Redundant PSUs (Power Supply Units)
2.13.4. Top Class AI-Ready Workstations for Studio-Level Performance
2.13.5. Consumer-Level AI-Friendly Solutions
2.13.6. Laptop Setups for Mobile Artists - The Hybrid Local + Cloud Workflow
2.14. Conclusion

3. Ethics & Responsibilities
3.1. Creative Integrity & Human Dignity
3.1.1. Authenticity vs Fabrication
3.1.1.1. Truthful Representation vs Creative Alteration
3.1.1.2. Addressing Deepfakes, Misinformation, and Harmful Manipulations
3.1.2. Consent and Identity
3.1.2.1. Image Likeness, Voice Replication, and Personal Data Usage
3.1.2.2. Securing Informed Consent from Living Persons and Estates of the Deceased
3.1.3. Labour and Creative Credit
3.1.3.1. Ensuring Human Contributions Are Acknowledged in AI-Assisted Work
3.1.3.2. Preserving Craft Recognition in Credits and Metadata
3.1.4. Ethics of Stylisation
3.1.4.1. Respecting Cultural Context When Adopting Artistic Styles
3.1.4.2. Avoiding Appropriation and Misrepresentation
3.1.5. Cultural Sensitivity & Cross-Border Considerations
3.1.5.1. Recognising Different Ethical Norms by Region and Audience
3.1.5.2. Adapting Creative Outputs for International Distribution and Reception
3.2. Rights, Law & Accountability
3.2.1. Ownership and Intellectual Property
3.2.1.1. Understanding Who Owns AI-Generated Works
3.2.1.2. Licensing Implications for Training Datasets and Outputs
3.2.2. Legal Framework vs Moral Compass: What is Right vs What Is Legal
3.2.2.1. IP Law Reference Texts
3.2.2.1.1. The Berne Convention
3.2.2.1.2. The Paris Convention
3.2.2.1.3. The TRIPS Agreement
3.2.2.1.4. WIPO Copyright Treaty (WCT)
3.2.2.1.5. Regional Frameworks
3.2.2.1.6. European Union – Copyright Directive (DSM Directive)
3.2.2.1.7. United States – Copyright Act & Copyright Office AI Guidance
3.2.2.1.8. United Kingdom – Copyright, Designs and Patents Act
3.2.2.1.9. Cross-Border Reality Check for AI-Created Works
3.2.2.2. Creative Commons vs Commercial Licensing of Models
3.2.2.2.1. Licensing Decision Tree for AI Models
3.2.2.3. Attribution in Community Models
3.2.2.4. Model Provenance and Metadata Tracking
3.2.2.5. Bias in Datasets
3.2.3. Model Accountability & Governance
3.2.3.1. Assigning Responsibility Between Model Creators, Dataset Curators, and End-Users
3.2.3.2. Liability in Cases of Misuse, Harm, or Rights Infringement
3.2.4. Evolving Legal Landscape / Staying Informed
3.2.4.1. Importance of Keeping Current With International AI and IP Laws
3.2.4.2. Following Legislative Developments and Industry Advocacy Groups
3.3. Responsible Practice & Sustainability
3.3.1. Bias and Representation
3.3.1.1. Avoiding Stereotypes and Reinforcing Inequalities
3.3.1.2. Balancing Creative Freedom With Social Responsibility
3.3.2. Data Privacy & Dataset Ethics
3.3.2.1. Handling Personal Data in Compliance with Laws
3.3.2.2. Best Practices for Anonymisation and Secure Storage
3.3.3. Transparency & Disclosure
3.3.3.1. Informing Audiences and Collaborators When AI Tools Are Used
3.3.3.2. Open Disclosure in Metadata, Production Notes, and Public Materials
3.3.4. Environmental Impact
3.3.4.1. Energy Costs of Training and Rendering AI Models
3.3.4.2. Carbon Footprint Reduction Strategies for AI Workflows
3.3.5. Misuse & Harm Mitigation
3.3.5.1. Proactively Identifying and Preventing Harmful Outputs
3.3.5.2. Creating Safeguards Against Unauthorised Model Use
3.3.6. Economic Sustainability & Industry Impact
3.3.6.1. Balancing Innovation with Preservation of Jobs and Craft Diversity
3.3.6.2. Ethical Adoption to Avoid Monopolisation by Large Tech Players
3.3.7. Continuous Education & Industry Advocacy
3.3.7.1. Lifelong Learning as an Ethical Obligation
3.3.7.2. Participating in Public Debate and Policy Shaping
3.3.8. Netflix's Creative AI Guidelines: A Real-World Reference for Generative Ethics
3.3.8.1. Generative AI Ethical Checklist for Artists

4. Using ComfyUI
4.1. Foundations
4.1.1. Introduction and Installation
4.1.1.1. What is ComfyUI?
4.1.1.2. Installing ComfyUI
4.1.1.2.1. System requirements
4.1.1.2.2. Portable vs Desktop Editions of ComfyUI
4.1.1.2.1. Folder Strategy: Centralized Models
4.1.1.2.2. Installation Steps
4.1.1.2.3. Basic Navigation inside ComfyUI
4.1.1.2.4. Downloading and Installing Models
4.1.1.3. First Workflow and First Image
4.1.1.3.1. Troubleshooting: What to do when workflow errors happen
4.1.1.4. Different Models, Different Needs: Working with SD 1.5 vs SDXL
4.1.1.4.1. Optimizing with Recommended Settings
4.1.1.5. Saving and Reusing Workflows
4.1.1.6. Using ComfyUI Manager
4.1.1.7. Closing and Shortcuts
4.1.1.8. Essential Hotkeys in ComfyUI
4.1.2. Nodes and Workflow Basics
4.1.2.1. Core Workflow: Text-to-Image
4.1.2.1.1. The ComfyUI Interface – Workflow Controls
4.1.2.2. The Canvas
4.1.2.3. Building a Text-to-Image Workflow from Scratch
4.1.3. Text-to-Image (TXT2IMG) Workflows
4.1.3.1. The KSampler Node and the Seed Parameter
4.1.3.2. Steps
4.1.3.3. CFG – Classifier-Free Guidance
4.1.3.4. Sampler and Scheduler
4.1.3.5. Denoise
4.1.3.6. Managing Jobs and Queues
4.1.3.7. Batch Generation
4.1.3.8. Output Management
4.1.3.9. Workflow Organization
4.1.3.10. Running Multiple Workflows
4.1.3.11. Best Parameters Ranges
4.1.4. Image-to-Image (IMG2IMG) Workflows
4.1.4.1. Denoising Strength in IMG2IMG
4.1.4.2. Image Resolution Considerations
4.1.4.3. Clipboard Workflow
4.1.4.4. Introduction to LoRAs
4.1.4.5. Downloading LoRAs
4.1.4.6. Using LoRAs in ComfyUI
4.1.4.7. Trigger Words and Strength
4.1.4.8. Combining LoRA with IMG2IMG
4.1.4.9. Recap of the Full Workflow
4.2. Models & Performance
4.2.1. Stable Diffusion 3 (SD3) Medium
4.2.1.1. Updating ComfyUI
4.2.1.2. Downloading SD3 Models
4.2.1.3. Loading SD3 in ComfyUI
4.2.1.4. Testing Multiple SD3 Variants
4.2.1.5. Comparing SD3 Variants
4.2.1.6. Comparing SD3 to SDXL
4.2.1.7. When to Use SD3 vs SDXL
4.2.2. Flux 1 Schnell & Dev – Install/Variants
4.2.2.1. Flux Versions and Licensing
4.2.2.2. Installing Flux Models
4.2.2.3. Building Workflows from Scratch
4.2.2.4. Flux Guidance Node
4.2.2.5. Styles Integration with Flux
4.2.2.6. Compact Workflow Version
4.2.2.7. Testing Flux Models (Comparison)
4.2.2.8. Installing the Full Dev Model (16-bit)
4.2.2.9. New NF4 Versions (Low VRAM)
4.2.2.10. LoRA with Flux
4.2.3. Flux GGUF & Custom Nodes

4.2.3.1. Managing Custom Nodes
4.2.3.2. Handling Missing Nodes in Workflows
4.2.3.3. Style Node Replacement
4.2.3.4. Custom Nodes Management in Practice
4.2.3.5. Installing Flux GGUF
4.2.3.6. Building a Flux GGUF Workflow
4.2.3.7. Comparing FP16 vs GGUF (Q8)
4.2.3.8. Compact Workflow Version
4.2.3.9. Testing Dev vs Schnell GGUF
4.2.3.10. Practical Prompt Tests
4.2.3.11. Searching for Custom Nodes
4.2.4. Fluxmania: Photorealistic Images
4.2.4.1. Installing Fluxmania
4.2.4.2. Building the Basic Workflow
4.2.4.3. Compact Workflow Variant
4.2.4.4. Prompting Strategies
4.2.4.5. Example Generations
4.2.4.6. Upscaling Workflow
4.2.4.7. Known Issues
4.3. Prompt Systems
4.3.1. LLM Integration (Florence/Surge)
4.3.1.1. Installing Florence (Image Captioning)
4.3.1.2. Using Florence for Image Captioning
4.3.1.3. Integrating Image Captioning into Workflows
4.3.1.4. Switching Between Prompts
4.3.1.5. Installing Surge LLM (Text-to-Text Prompt Generation)
4.3.1.6. Using Surge LLM for Prompt Generation
4.3.1.7. Integrating Text-Based LLM into Workflows
4.3.1.8. Experiments and Limitations
4.3.1.9. Compact Prompt Generation Workflows
4.3.2. LLMs via Ollama (Gemma/LLaVA)
4.3.2.1. Installing Ollama
4.3.2.2. Exploring Ollama Models
4.3.2.3. Integrating Ollama into ComfyUI
4.3.2.4. Understanding Generation Settings
4.3.2.4.1. Core knobs
4.3.2.4.2. Recommended presets (copy the vibe, not dogma)
4.3.2.5. Testing Gemma Models
4.3.2.6. Using Vision Models (LLaVA)
4.3.2.7. Building Hybrid Workflows
4.3.2.8. Practical Use Cases
4.3.3. Prompt Styles / Prompt Loader / Batch
4.3.3.1. The Styles File
4.3.3.1.1. Styles Library – Documentation
4.3.3.2. Installing the I-Tools Custom Node
4.3.3.3. Where the Style Templates Live (and How to Enable More)
4.3.3.4. One-Style vs Multi-Style Nodes
4.3.3.5. Wire Styles into Your Workflows (SDXL & Flux)
4.3.3.6. Editing or Authoring Your Own Style Files (YAML)
4.3.3.7. Text Overlays on Images (Watermarks, Prompts, Identifiers)
4.3.3.8. Batch: Load Many Images from a Folder (+ Optional Crop)
4.3.3.9. Use the Image Filename as the Prompt (Batch i2i / Upscale)
4.3.3.10. Run Multiple Prompts from a File (Prompt Loader)
4.3.3.11. Quick SDXL & Flux Wiring Recap
4.4. Structural Control
4.4.1. SDXL ControlNet Union
4.4.1.1. Installing the ControlNet Model
4.4.1.2. Understanding ControlNet
4.4.1.3. Updating ComfyUI Interface
4.4.1.4. Installing Custom Nodes for ControlNet
4.4.1.5. Working with Pre-Processors
4.4.1.6. Integrating ControlNet into an SDXL Workflow
4.4.1.7. Strength Control
4.4.1.8. Image Quality and Ratios
4.4.1.9. Stacking Multiple ControlNets
4.4.1.10. Practical Examples
4.4.2. Flux ControlNet Union Pro
4.4.2.1. Installing the ControlNet Model
4.4.2.2. Required Custom Nodes
4.4.2.3. Inserting ControlNet into a Flux Workflow
4.4.2.4. Preprocessors & Image Preparation
4.4.2.5. Workflow Assembly
4.4.2.6. Flux Resolution Calculator
4.4.2.7. Practical Examples
4.4.2.8. Dev vs Schnell
4.4.3. ControlNet Union Pro 2.0
4.4.3.1. Setup Requirements
4.4.3.2. Workflow 1: Canny Edge ControlNet
4.4.3.3. Workflow 2: Depth Anything v2 Pre-Processor
4.4.3.4. Workflow 3: Converting to Image-to-Image
4.4.3.5. Workflow 4: Pose Pre-Processor (DW Recommended)
4.4.3.6. Workflow 5: Any Line Art Pre-Processor
4.4.3.7. Workflow 6: Combining Multiple Pre-Processors
4.4.3.8. Workflow 7: Using LoRAs with ControlNet
4.4.3.9. Workflow 8: Fluxmania + ControlNet
4.4.3.10. Troubleshooting & Tips
4.5. Editing & Paint-Over Tools
4.5.1. Inpainting
4.5.1.1. Installing Inpainting Nodes
4.5.1.2. Setting Up the Workflow (SDXL)
4.5.1.3. Mask Editor Basics
4.5.1.4. Using Inpainting for Edits
4.5.1.5. Context Area & Mask Options
4.5.1.6. Practical Examples
4.5.1.7. Outpainting (SDXL)
4.5.1.8. Flux Inpainting & Outpainting
4.5.1.9. Photoshop as a Companion Tool
4.5.2. Remove Backgrounds
4.5.2.1. Installing InSPyReNet Rem BG Node
4.5.2.2. Basic Background Removal Workflow in ComfyUI
4.5.2.3. Saving Masks for Editing
4.5.2.4. Example Results and Limitations
4.5.2.5. Using the Advanced InSPyReNet Node
4.5.2.6. Photoshop for Refinement
4.5.2.7. ComfyUI + Photoshop Hybrid Workflow
4.5.2.8. ComfyUI + Nuke Hybrid Workflow
4.5.2.9. Practical Example: Clip Art / Game Assets
4.5.3. Backgrounds Replacement
4.5.3.1. Workflow Setup (Replace Complex Background)
4.5.3.2. Working with Alpha Masks
4.5.3.3. Prompting for Backgrounds
4.5.3.4. Troubleshooting Artifacts
4.5.3.5. Special Case: White / Solid Backgrounds
4.5.3.6. Fine-Tuning Results
4.5.3.7. People & Hair Edges
4.5.3.8. Changing the Subject Instead of the Background
4.5.3.9. Image Ratios & Limitations
4.5.4. Flux Tools
4.5.4.1. Installation & Setup
4.5.4.2. Flux Fill (Inpainting & Outpainting)
4.5.4.3. Flux Redux (Variations & Multi-Image Fusion)
4.5.4.4. Flux Depth
4.5.4.5. Flux Canny
4.5.4.6. LoRA Versions (Depth & Canny)
4.5.4.7. Practical Notes
4.5.5. Flux Redux & LoRA Inpainting
4.5.5.1. Installing Advanced ReFlux Control Node
4.5.5.2. Controlling Image Strength in Flux Redux
4.5.5.3. Redux Advanced Node with Masking
4.5.5.4. Multi-Image Redux Workflow
4.5.5.5. Using LoRA for Inpainting with Flux
4.5.5.6. Training and Using Custom LoRAs
4.5.5.7. Best Practices for LoRA Inpainting
4.5.6. Flux Kontext for Editing & Consistency
4.5.6.1. Installation & Setup
4.5.6.2. Basic Image Generation with Flux Kontext
4.5.6.3. Upscaling Workflow
4.5.6.4. Editing & Transformation Workflows
4.5.6.5. Combining Images
4.5.6.6. Style Transfer & Transformation

4.5.6.7. Branding & Product Placement
4.5.6.8. Color, Text & Environmental Edits
4.5.6.9. Character Replacement & Consistency
4.5.6.10. Background Replacement
4.5.6.11. Object Removal & Inpainting
4.5.6.12. LoRA Integration
4.5.6.13. Multi-Character Workflows
4.5.7. Cloud Usage
4.6. Identity & Style Systems
4.6.1. LoRAs in Flux
4.6.1.1. Preparing ComfyUI & New Menu UI
4.6.1.2. Adding LoRAs to a Workflow
4.6.1.3. Downloading and Organizing LoRAs
4.6.1.4. Quick SDXL & Flux Wiring Recap
4.6.1.5. Using LoRAs in Practice
4.6.1.6. Examples of LoRA Types
4.6.1.7. Combining LoRAs
4.6.1.8. LoRAs with Image-to-Image
4.6.1.9. Performance Considerations
4.6.1.10. Community & Resources
4.6.2. Character LoRA Pipeline
4.6.2.1. Character Sheet Generation
4.6.2.2. Using OpenPose for Pose Control
4.6.2.3. Character Sheet Results & Seeds
4.6.2.4. Preparing Images for Training
4.6.2.5. Alternative Characters (Boy, Alien, Animals)
4.6.2.6. Human Characters (No ControlNet)
4.6.2.7. Creating More Variations from Limited Data
4.6.2.8. Training LoRAs Online
4.6.2.9. Using Trained LoRAs in ComfyUI
4.6.2.10. Upscaling & Variations with LoRAs
4.6.2.11. Expanding Beyond Images (Image-to-Video)
4.6.3. Live Portrait
4.6.3.1. Installing Required Nodes
4.6.3.2. Basic Workflow with Expression Editor
4.6.3.3. Crop Factor Considerations
4.6.3.4. Advanced Live Portrait with Driving Video
4.6.3.5. Combining Expression Editor with Driving Video
4.6.3.6. Animate Without Video
4.6.3.7. Crop and Blending Issues
4.6.3.8. Limitations & Best Practices
4.6.4. Flux Kontext LoRA Training (Fal AI)
4.6.4.1. Why Fal AI?
4.6.4.2. Preparing Training Data
4.6.4.3. Training on Fal AI
4.6.4.4. Using the LoRA in Flux Kontext
4.6.4.5. Converting LoRA for Nunchaku
4.6.4.1. Practical Issues & Fixes
4.6.4.2. Beyond Icons – Training Characters
4.6.4.1. Experimental LoRAs
4.6.5. USO: Unified Style and Subject-Driven Generation
4.6.5.1. What is USO Actually Doing?
4.6.5.2. Choosing Subject and Style References
4.6.5.3. The USO Prompt Structure: Shaping the Outcome
4.6.5.4. Setting Up the USO Workflow in ComfyUI
4.6.5.5. Tuning Parameters for Subject and Style Weight
4.6.5.6. Style Transfer Use Cases
4.6.5.7. Bonus Tips and Troubleshooting
4.7. Creative Modalities
4.7.1. Sketch-to-Image
4.7.1.1. Installing and Updating Nodes
4.7.1.2. Overview of the All-in-One Workflow
4.7.1.3. Sketch Preparation Workflow
4.7.1.4. Image Variation Workflow
4.7.1.5. Sketch to Image Workflow
4.7.1.6. Draw Your Sketch to Image Workflow
4.7.1.7. Photoshop as a Companion Tool
4.7.1.8. SDXL vs. Flux Workflows
4.7.1.9. Beyond Still Images
4.7.2. Photo-to-Cartoon Transformation Workflow
4.7.2.1. Choosing Cartoon Models & LoRAs
4.7.2.2. Prompt Structuring for Cartoons
4.7.2.3. Denoise Strength for Cartoonization
4.7.2.4. ControlNet for Structure Preservation
4.7.2.5. Hybrid Workflow: SDXL + Flux
4.7.2.6. Special Considerations for Non-Character Subjects
4.7.3. OmniGen
4.7.3.1. Installing OmniGen in ComfyUI
4.7.3.2. First Run and Model Download
4.7.3.3. Using the OmniGen Node
4.7.3.4. Troubleshooting Errors
4.7.3.5. Multi-Image Workflows
4.7.3.1. Practical Use Cases
4.7.3.2. Performance Notes
4.7.3.3. Integration with Other Workflows
4.7.4. Photo to Watercolor/Oil
4.7.4.1. Required Models & Setup
4.7.4.2. Workflow Structure
4.7.4.3. Watercolor Workflow Example
4.7.4.4. Upscaling Stage
4.7.4.5. Oil Painting Workflow Example
4.7.4.6. Landscape Example
4.7.4.7. Impressionist Oil Workflow Example
4.7.4.8. Digital Painting Workflow Example
4.7.5. DeepSeek Janus Pro (Image-to-Prompt)
4.7.5.1. Installing Janus Pro
4.7.5.2. Downloading the Model Weights
4.7.5.3. Loading the Model in ComfyUI
4.7.5.4. Image Understanding Workflow
4.7.5.5. Using Janus Prompts with Generators
4.7.5.6. Enhancing Prompts with Text Concatenation
4.7.5.7. Testing on Different Images
4.7.5.8. Improving Prompt Instructions
4.7.5.9. Comparing 1B vs 7B Models
4.7.5.10. Integration with Other Workflows
4.7.6. Qwen Wrokflows I: Image Generation
4.7.6.1. Setup & Updates
4.7.6.2. Workflow Basics
4.7.6.3. Performance & Resolutions
4.7.6.4. Strengths
4.7.6.5. Weaknesses
4.7.6.6. Image-to-Image Workflow
4.7.6.7. Comparisons
4.7.6.8. Civic AI Nunchaku Model
4.7.7. Qwen Workflows II: Image Editing
4.7.7.1. Setup & Installation
4.7.7.2. Core Workflow
4.7.7.3. Example Workflows
4.7.7.4. Editing Categories
4.7.7.5. Logos & Products
4.7.7.6. Clothes & Try-On Workflows
4.7.7.7. Observations & Limitations
4.7.8. Seamless Pattern & Texture Generation
4.7.8.1. Preparing ComfyUI
4.7.8.2. Building the Base Workflow
4.7.8.3. First Results & Photoshop Test
4.7.8.4. Checking Seamlessness Inside ComfyUI
4.7.8.5. Tiling Enabled vs Disabled
4.7.8.6. Prompt Engineering for Patterns
4.7.8.7. Offset Image Node
4.7.8.8. Output Organization
4.7.8.9. Applications of Patterns
4.7.8.10. Limitations
4.8. Quality & Delivery
4.8.1. Upscaling Fundamentals
4.8.1.1. Installing Upscaler Nodes and Models
4.8.1.2. Basic Upscale Workflow
4.8.1.3. Choosing the Right Upscale Model
4.8.1.4. Creative Upscaling with Flux
4.8.1.5. Image-to-Image Upscaling with KSampler

4.8.1.6. Multi-Step Upscaling
4.8.1.7. Comparing Flux vs SDXL Upscaling
4.8.1.8. Practical Examples
4.8.1.9. Compact Upscale Workflows
4.8.2. Advanced Upscaling Workflows
4.8.2.1. SDXL Upscale Workflow
4.8.2.2. Flux Dev Upscale Workflow
4.8.2.3. Fluxmania Upscale Workflow
4.8.2.4. SD v1.5 Upscale Workflow (Low VRAM Option)
4.8.2.5. Text-to-Image + Upscale Workflows
4.8.2.6. Cloud Upscaling (Running Hub)
4.8.2.7. Example Results & Comparisons
4.9. Motion
4.9.1. LTX 0.9.7 Animation Workflows
4.9.1.1. Installation & Setup
4.9.1.2. Base Workflow: Image-to-Video
4.9.1.3. Upscaling Workflow
4.9.1.4. Multi-Frame Input Workflow
4.9.1.5. Aspect Ratios & Formats
4.9.1.6. Extending Videos
4.9.1.7. Low-VRAM (GGUF) Workflow
4.9.1.8. Example Project: Alien Short Film
4.9.2. Wan 2.1 VACE Workflows
4.9.2.1. Installation & Setup
4.9.2.2. Workflow 1: Text-to-Video
4.9.2.3. Workflow 2: Image-to-Video
4.9.2.4. Workflow 3: Video-to-Video
4.9.2.5. Optimization with LoRA (Speed & Quality)
4.9.2.6. Fixing Contrast Issues
4.9.2.7. Performance & Comparisons
4.10. Workflow Power-Ups
4.10.1. Customizing Model Paths in ComfyUI
4.10.2. Tips, Tricks & Best Practices
4.10.2.1. Updating ComfyUI (Manager vs Manual Updates)
4.10.2.2. Interface Customization (New Menu, Resource Monitor, Workflow Tabs)
4.10.2.3. Node Management (Badges, Link Styles, Groups & Muting, Renaming)
4.10.2.4. Workflow Saving, Exporting, and Bookmarks
4.10.2.5. Save Image Node Tricks (Prefixes, Dates, Folders)
4.10.2.6. Previews & History (Live Preview, Queue, Reloading Workflows from Images)
4.10.2.7. VRAM Management (VAE Decode Tiled, Unload Models)
4.10.2.8. Notes, Labels, and Collapsing Nodes
4.10.2.9. Copy/Paste Shortcuts (Keeping Links Intact)
4.10.2.10. Metadata in PNGs (Loading Workflows from Images, Missing Nodes)
4.10.2.11. IT Tools & Multi-Prompt Workflows (Increment, Queue Instant)
4.10.2.12. Debugging & Error Handling (Logs, GitHub Issues, Discord FAQ)
4.10.2.13. Node Pinning, Themes, and Custom Node Settings
4.10.2.14. Embedded Python Environment (Checking Dependencies, Running Scripts)
4.10.3. Workflow Utilities with KayTool
4.10.3.1. Installing KayTool
4.10.3.2. Resource Monitor
4.10.3.3. Workflow Export as PNG
4.10.3.4. Node Alignment Tools
4.10.3.5. Wireless Connections (Set/Get Nodes)
4.10.3.6. Slider Controls
4.10.3.7. Display Any Node
4.10.3.8. Color Adjustment Node
4.10.3.9. Image Size Nodes
4.10.4. Documentation & GitHub
4.10.5. Nunchaku for Flux Dev & Kontext
4.10.5.1. Installation with Easy Installer
4.10.5.2. Linking Models from Existing ComfyUI
4.10.5.3. Running Nunchaku
4.10.5.4. Setting Up Nunchaku Workflows
4.10.5.5. Performance Gains
4.10.5.6. Fixing Common Errors
4.10.5.7. Adapting Old Workflows
4.10.5.8. Example Workflows
4.10.5.9. Observations
4.10.6. Wan 2.2 for Image Generation
4.10.6.1. Setup & Requirements
4.10.6.2. Workflow Basics
4.10.6.3. Upscaling
4.10.6.4. Example Prompts & Results
4.10.6.5. Image-to-Image Workflow
4.10.6.6. Social Media Workflows
4.10.6.7. LoRA Integration
4.10.7. Workflow Optimization with TeaCache
4.10.7.1. Installing TeaCache
4.10.7.2. First Example – Flux Dev Workflow
4.10.7.3. Second Example – WAN + LoRA Workflow
4.10.7.4. How TeaCache Works
4.10.7.5. Key Parameters
4.10.8. Running ComfyUI in the Cloud
4.10.8.1. Platform Overview – RunningHub
4.10.8.2. Interface Tour
4.10.8.3. Running Workflows in the Cloud
4.10.8.4. Importing Your Own Workflows
4.10.8.5. Managing Generations
4.10.8.6. Model Management
4.10.8.7. Example – ControlNet Workflow
4.10.8.8. Subscriptions and Costs

5. ComfyUI + Nuke: Generative AI Workflows for Compositing

5.1. Introduction: The New Frontier of Compositing
5.1.1. How Generative AI expands Nuke's creative potential
5.1.2. The philosophy of integration: when node logic meets neural networks
5.1.3. Why AI belongs in the compositor's toolkit, not outside it
5.2. Understanding the Integration
5.2.1. What ComfyUI is and how it parallels Nuke's node paradigm
5.2.2. The architecture of the bridge: Nuke as client, ComfyUI as server
5.2.3. JSON workflows, WebSocket communication, and model handling
5.2.4. The concept of "generative nodes" inside Nuke
5.3. Installing and Configuring the ComfyUI Bridge
5.3.1. Preparing the environment (Python, Git, ComfyUI)
5.3.2. Installing ComfyUI-for-Nuke step-by-step
5.3.3. Enabling EXR support and color-managed output (HQ Image Save node)
5.3.4. Folder structure, port configuration, and connection tests
5.3.5. Troubleshooting common setup issues
5.4. Navigating the Node Ecosystem
5.4.1. How Nuke reads ComfyUI nodes and generates its counterparts
5.4.2. Using the ComfyUI menu inside Nuke
5.4.3. Importing and editing workflows (.json)
5.4.4. Understanding node naming and parameter translation
5.4.5. Working efficiently with node graphs (caching, lazy loading, model management)
5.5. Running AI Generations from Within Nuke
5.5.1. The QPrompt node: sending text prompts to ComfyUI
5.5.2. Executing image generations directly from the Nuke graph
5.5.3. Managing seed, steps, CFG, and other key diffusion parameters
5.5.4. Understanding render latency and optimizing inference speed
5.5.5. Batch generation and remote GPU configurations
5.6. ComfyUI Workflows for VFX Tasks
5.6.1. Matte Painting and Set Extensions
5.6.1.1. From text prompt to photoreal matte painting
5.6.1.2. Integrating AI backgrounds behind live-action plates
5.6.1.3. Using ControlNet for horizon and composition alignment
5.6.1.4. Matching lighting and color grading for seamless blends
5.6.2. Depth and Projection Workflows
5.6.2.1. Generating AI depth maps from 2D images
5.6.2.2. Creating 2.5D projections and parallax camera moves
5.6.2.3. Combining MiDaS/ZoeDepth with Nuke's 3D system
5.6.2.3.1. Using AI normals for relighting and fog passes
5.6.3. Roto, Masking, and Cleanup
5.6.3.1. The MatAnyOne model for automated video mattes
5.6.3.2. Combining SAM and generative inpainting for object removal

5.6.3.3. Generating clean plates with DiffuEraser
5.6.3.4. Quality checks and limitations (temporal stability, edge artifacts)
5.6.4. Lighting, Texture, and Lookdev
5.6.4.1. Generating HDRIs and environment maps via AI
5.6.4.2. Texturing 3D renders with depth/normal ControlNets
5.6.4.3. Using AI for relighting and texture detailing
5.6.4.4. Integrating generated maps into Nuke's lighting tools
5.6.5. Upscaling and Detail Enhancement
5.6.5.1. ESRGAN, SwinIR, and diffusion-based super-resolution
5.6.5.2. Preserving filmic grain and fine detail
5.6.5.3. Regional upscaling and facial enhancement use cases
5.7. Practical Production Workflows
5.7.1. AI-assisted previsualization and storyboarding
5.7.2. Using LoRAs for consistent visual style across shots
5.7.3. Shot-based iterations: from concept to comp-ready asset
5.7.4. Round-trip workflows: sending Nuke renders into ComfyUI and back
5.7.5. Multi-artist and server setups in studio environments
5.8. Creative & Experimental Techniques
5.8.1. ControlNet Magic
5.8.1.1. Depth, pose, scribble, and normal conditioning
5.8.1.2. From RotoPaint sketches to fully realized images
5.8.1.3. Multi-ControlNet setups for ultimate creative control
5.8.2. WAN and Flux Models
5.8.2.1. Using WAN 2.x for video interpolation, morphing, and face swaps
5.8.2.2. Flux 1: higher fidelity image generation for production use
5.8.2.3. Comparative study: SDXL vs Flux vs WAN
5.8.2.4. Ethical and technical considerations for generative video
5.8.3. Hybrid Pipelines and Cross-Software Collaboration
5.8.3.1. Sending renders from Houdini/Blender into Nuke-ComfyUI
5.8.3.2. Building a real-time "AI Comp Lab" with networked workstations
5.8.3.3. Future outlook: real-time diffusion and GPU sharing
5.9. Building Custom AI Gizmos
5.9.1. Encapsulating AI workflows as artist-friendly nodes
5.9.2. Example: "Storyboard Generator" and "Quick Sky Replacement" gizmos
5.9.2.1. Example Gizmo I: "Storyboard Generator" (inside Nuke)
5.9.2.2. Example Gizmo II: "Quick Sky Replacement" (inside Nuke)
5.9.3. Exposing the right controls: prompts, seeds, and model presets
5.9.4. Best practices for tool design, naming, and maintenance
5.10. Troubleshooting and Optimization
5.10.1. Handling model mismatches and missing nodes
5.10.2. Debugging WebSocket and JSON errors
5.10.3. GPU memory management and multi-GPU setups
5.10.4. Tips for caching, previewing, and error recovery
5.11. Beyond Integration: The Future of AI Compositing
5.11.1. How AI transforms the role of the compositor
5.11.2. Ethical use of AI in film and VFX
5.11.3. From experimental tools to industry standards
5.11.4. What's next: SD 3.0, NeRF, and real-time generative rendering

6. Community Tools & Resources

6.1. GIT Repositories
6.1.1. Understanding Git and GitHub
6.1.2. Official & Core Repositories
6.1.3. High-Impact Community Node Packs
6.1.4. Specialised Utility Nodes
6.1.5. ComfyUI-Nuke API Repository
6.2. AI Community Platforms
6.2.1. Website-Based Platforms
6.2.2. ComfyUI Discord Server
6.2.3. On Reddit
6.3. Image Models Arenas
6.3.1. How to Use Image Arenas to Pick the Right Model for ComfyUI
6.4. Utilities
6.5. Appendices (PDFs)
6.5.1. Installation Summary
6.5.1.1. Windows Installation
6.5.1.2. Installing ComfyUI-Custom-Scripts by PythonGosssss (Windows, Official Method)
6.5.1.3. Mac Installation

Index

Note: The curriculum and structure outlined in this syllabus are subject to refinement or modification prior to and during the course delivery, to ensure optimal relevance and pedagogical clarity.



Victor Perez

Victor Perez is a renowned figure in the global visual effects industry. His career spans over 28 years so far, and he has worked in many areas of visual storytelling. Starting as a photographer and cinematographer, extending his interest to the whole image post-production process and evolving into a visual effects artist with a special focus on photorealistic compositing of CGI. Along the way, he has worked on a number of Hollywood films, including Christopher Nolan's 'The Dark Knight Rises', 'Rogue One: A Star Wars Story', 'Harry Potter and the Deathly Hallows', and 'Pirates of the Caribbean: On Stranger Tides', plus many more. Eventually, he became a Film Director, and his career naturally led him to his current role as a Visual Effects Supervisor.

What really drives Perez is his focus on storytelling. He is fascinated by how new technologies can help filmmakers tell stories in ways never seen before. His work is a fine alchemy of Science, Technology, and Art, reflecting his passion for innovation and creativity.

In parallel and compatible with his professional career as a production visual effects supervisor, Victor Perez has built a name for himself in the visual effects industry worldwide for the last 15 years by sharing his knowledge and work with the global community of visual effects professionals. Regularly invited to renowned events such as Siggraph, FMX, VIEW Conference, and more, he has become an influential figure in the field. His insights and accomplishments have been featured in specialised press, including Cinefex, FXGuide, The VFX Voice, and Before & Afters, to mention a few. In 2012, he was voted one of the most influential visual effects artists in the world by the specialised press of Nukepedia.

Acting as Brand Ambassador for ASUS, Autodesk, Eizo, Wacom and Foundry, Perez receives sponsorship, the latest technology, and financial support for his pioneering artistic research. Internationally regarded as a VFX compositing guru, he has delivered Masterclasses on Photorealistic Full-CGI Compositing at industry-leading institutions like Pixar Animation Studios and Dneg; and educational institutes like Escape Studios, and The Animation Workshops.

In 2016, he was invited to speak at TED, showcasing his research and its connection with his personal life, featuring a live demo of the visual effects compositing process. Perez's publications on various facets of visual effects are available through Routhledge, Netflix, FxPhD, and other outlets. Author of the industry-standard best selling book "The Color Management Handbook for Visual Effects Artists" published by Focal Press.

Victor Perez's multifaceted career in filmmaking has been honoured with over 50 awards and nominations. He has been recognised for his roles as visual effects supervisor, director, writer, and producer. Among these honours there are two prestigious Italian Academy Awards "David di Donatello," which he received, in 2019 for his outstanding work on "The Invisible Boy: Second Generation", and in 2025 for his internationally acclaimed work on "Naples – New York." In 2019, he made history with his experimental project "Echo," which he also directed. It earned a nomination at the Visual Effects Society (VES) Awards for Best Virtual Cinematography, marking the first time a short film had ever been nominated for this prestigious award.

Through his extensive career, Victor Perez continues to be a beacon of innovation and divulgation in the visual effects community, bridging the gap between technology, creativity, and storytelling.

More Information and Orders
www.victorperez.co.uk

Contact Us
info@victorperez.co.uk
+44 (0) 7757 199-766

Masked Frame Pictures Ltd
3rd Floor
207 Regent Street
London W1B 3HH
United Kingdom

We Offer Volume Licensing for Companies and Educational Institutions
info@victorperez.co.uk

© 2025 Victor Perez. All rights reserved.
This program and its contents are the intellectual property of Victor Perez. No part of this material may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the copyright holder.
Victor Perez™ is a registered trademark and artistic signature of Victor Raul Perez Raya. Unauthorised use of this trademark is strictly prohibited.
The curriculum is subject to changes and updates prior to and during the course delivery to reflect evolving technologies, tools, and industry best practices.