

Supplementary Material

DuoGen: Towards Autonomous Interleaved Multimodal Generation

Supplementary Materials

1. Interleaved Generation Examples	2
2. Image Generation and Editing Examples	10
3. Comparison and Results of More VLM Judge	13
4. Detailed Results on Image Generation and Editing Benchmarks	14
5. Details of the Interleaved Generation Benchmark	15
6. Data Engine for Webpages	17
7. Synthetic Interleaved Generation Data	22
8. Interleaved Context Alignment Data	24

1. Interleaved Generation Examples


Fig. A1 and Fig. A2 present additional examples from the Cooking-200. Fig. A3, Fig. A4, Fig. A5, and Fig. A6 show examples from the How-to-500. Fig. A7 and Fig. A8 provide results on InterleavedBench [18].

 Can you teach me how to make medium-rare steak?



1

 Can you teach me how to make Fried Rice?



2

 Can you teach me how to make Mushroom and Egg Stir-Fry?



3



****Step 1: Prepare Your Steak****
Start with a good cut of steak, such as a ribeye or New York strip. Pat it dry thoroughly with paper towels. This is crucial for getting a good sear. Season generously with salt and freshly ground black pepper on all sides.



****Step 2: Preheat Your Pan****
Place a heavy-bottomed pan, preferably cast iron, over high heat. Let it heat up for several minutes until it's smoking slightly. Add a tablespoon of high smoke point oil, like canola or grapeseed oil, to the pan.



****Step 3: Sear the Steak****
Carefully place the seasoned steak in the hot pan. Sear undisturbed for 2-3 minutes per side to develop a beautiful crust. You should hear a good sizzle.



****Step 4: Baste the Steak****
Once both sides are seared, reduce the heat to medium. Add a knob of butter, a few cloves of crushed garlic, and some fresh rosemary or thyme sprigs to the pan. Tilt the pan slightly and, using a spoon, baste the melting butter and aromatics over the steak for 2-3 minutes, flipping the steak once or twice during this process.



****Step 5: Rest the Steak****
Transfer the cooked steak to a cutting board. It's crucial to let the steak rest for at least 5-10 minutes (for a smaller cut) or up to 15-20 minutes (for a larger cut). This allows the juices to redistribute throughout the meat, resulting in a more tender and flavorful steak.



****Step 6: Slice and Serve****
Once rested, slice the steak against the grain into desired thickness. Arrange the slices on a plate. You can drizzle any pan juices over the top or add a final sprinkle of fresh herbs if desired. Enjoy your perfectly medium-rare steak!

1



Let's get cooking! Here are the steps to make delicious fried rice, just like the image you provided.
****Step 1: Prepare Your Ingredients****
Gather all your ingredients. For this fried rice, you'll need cooked day-old rice, your choice of protein (like chicken, shrimp, or tofu), an egg, and vegetables such as peas, carrots, and onions. Chop everything into small, bite-sized pieces.



****Step 2: Cook Your Protein****
Heat a large wok or skillet over medium-high heat. Add a little oil, then cook your chosen protein until it's nearly done. Remove it from the wok and set it aside.



****Step 3: Sauté the Vegetables****
Add a bit more oil to the hot wok. Add the onions and carrots first, stir-frying for a couple of minutes until slightly tender. Then add the peas and cook for another minute.



****Step 4: Add the Rice and Egg****
Push the vegetables to one side of the wok. Pour in the day-old cooked rice, breaking up any clumps. Spread it out evenly. Make a small well in the center and crack an egg into it. Scramble the egg until fully cooked, then mix it into the rice and vegetables.



****Step 5: Combine and Season****
Return the cooked protein to the wok. Add soy sauce and a dash of sesame oil (optional) and stir everything together. Continue to stir-fry for a few more minutes, ensuring everything is heated through and well combined.



****Step 6: Serve and Enjoy!****
Transfer your delicious homemade fried rice into a bowl. It's ready to be enjoyed!

2




Here's how to make that delicious Mushroom and Egg Stir-Fry!
****Step 1: Prepare the ingredients****
Slice your mushrooms (shiitake or similar) and vegetables (chilies, green peppers, garlic) as shown. Beat the eggs in a bowl.



****Step 2: Cook the eggs****
Heat a wok or large pan with a little oil. Pour in the beaten eggs and scramble them until just set but still soft. Remove the cooked eggs and set aside.



****Step 3: Stir-fry the mushrooms and garlic****
Add a little more oil to the wok. Add the sliced mushrooms and stir-fry until they start to soften and release their moisture.



****Step 4: Add the remaining vegetables****
Add the green peppers and chili slices to the wok with the mushrooms. Continue to stir-fry for a few minutes until the vegetables are tender-crisp.



****Step 5: Combine and season****
Return the cooked eggs to the wok with the mushrooms and vegetables. Add your preferred stir-fry sauce ingredients (e.g., soy sauce, oyster sauce, a pinch of sugar). Toss everything together to combine and heat through.




****Step 6: Serve and enjoy!****
Transfer the stir-fry to a serving dish.

3


Figure A1. Generation results of DuoGen on Cooking-200.

Can you teach me how to make Mushroom and Egg Stir-Fry?



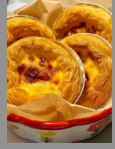
1

Can you teach me how to make borscht?





2


Can you teach me how to make egg tarts?





3














Here's a step-by-step guide to making delicious steamed fish, complete with images!

****Step 1: Prepare the Fish****
First, ensure your fish is fresh and thoroughly cleaned. Make sure it's scaled and gutted. You can score the fish diagonally on both sides to help it cook evenly and absorb flavors. Pat it dry with paper towels.

****Step 2: Add Aromatics****
Stuff the cavity of the fish with sliced ginger and green onions. You can also place some on top and under the fish. These aromatics will infuse the fish with wonderful flavor as it steams.


****Step 3: Prepare the Steamer****
Fill a wok or a large pot with water and bring it to a boil. Place a steaming rack inside.


****Step 4: Steam the Fish****
Carefully place the plate with the fish onto the steaming rack. Cover the wok/pot and steam for about 8-10 minutes, or until the fish is cooked through. The cooking time will vary depending on the size and thickness of your fish. You'll know it's done when the flesh is opaque and flakes easily with a fork.


****Step 5: Prepare the Sauce (Optional)****
While the fish is steaming, you can prepare a simple sauce. Combine soy sauce, a touch of sesame oil, and a pinch of sugar. You can also add some freshly ground black pepper.


****Step 6: Finishing Touches****
Once the fish is cooked, carefully remove it from the steamer. Discard the old ginger and green onions. You can add fresh julienned ginger and green onions on top for garnish.


1














Here's a step-by-step guide to making borscht, complete with images!

First, gather your ingredients. You'll need beets, cabbage, potatoes, carrots, onions, beef (or vegetable broth for a vegetarian version), tomato paste, and dill.

Next, chop all your vegetables. Dice the onions, carrots, and potatoes, shred the cabbage, and grate the beets.


In a large pot, sauté the onions and carrots until softened.


Then, add the grated beets and cook for a few more minutes.


Now, add the diced potatoes, cabbage, beef broth, tomato paste, and seasonings. Bring to a boil, then reduce heat and simmer until the vegetables are tender.


Finally, ladle the hot borscht into a bowl and garnish with fresh dill and a dollop of sour cream. Enjoy your delicious homemade borscht!


2














Let's make some delicious egg tarts!

First, you'll want to prepare your pastry dough. This typically involves combining flour, butter, and a bit of water to form a pliable dough. After chilling, roll it out thinly.

Next, cut out circles from your rolled-out dough and press them into individual tart molds, forming the crust.

While the crust is chilling in the molds, prepare the egg custard filling. This usually involves whisking together eggs, sugar, evaporated milk, and sometimes a touch of vanilla extract. Make sure the mixture is smooth.

Pour the egg custard mixture into the pastry-lined molds, filling them almost to the top. Be careful not to overfill!

Then, carefully place the molds into a preheated oven. Bake them until the crust is golden brown and the custard is set but still slightly wobbly in the center.

When cooled, you'll have beautiful, flaky egg tarts ready to enjoy! You can serve them warm out of the oven or chilled, perhaps with a dusting of powdered sugar or a drizzle of honey.

3

Figure A2. Generation results of DuoGen on Cooking-200.

 Can you teach me how to craft a fabric tassel keychain as a small gift?



Step 1: Gather Your Materials Before you begin, ensure you have all the necessary materials. You will need: fabric scraps (cotton, linen, or even an old t-shirt works well), a sharp pair of scissors, a ruler, a sewing machine (or needle and thread), an iron, a jump ring, and a keychain ring.


Step 2: Prepare Your Fabric Strips Cut your fabric into strips. For a standard tassel, I recommend strips that are about 1.5 to 2 inches wide and 10-12 inches long. You'll need about 4-6 strips for a nice full tassel. Iron them flat for easier handling.

Step 3: Fold and Press Take each strip and fold it in half lengthwise, pressing it firmly with your iron. This will give you a crisp edge.

Step 4: Stack and Stitch Stack your folded fabric strips neatly on top of each other. Then, using your sewing machine or needle and thread, sew a straight line across the raw edges, about 1/4 inch from the top. Backstitch at the beginning and end to secure the stitches.

Step 5: Create the Loop Take one of the fabric strips and create a small loop at the top. This loop will be used to attach the jump ring and keychain. You can secure this with a few hand stitches or a dab of fabric glue.

Step 6: Finish Your Tasse Now, carefully cut slits into the folded edges of your fabric strips, stopping just before your stitched line. Make these slits about 1/4 to 1/2 inch apart to create the fringe of your tassel. Once you're done, attach a jump ring through your fabric loop, and then thread your keychain ring onto the jump ring. And there you have it 2013 a beautiful, handmade fabric tassel keychain

 Can you teach me how to craft a sturdy tomato cage trellis from wire mesh?



Step 1: Gather Your Materials Before you begin, ensure you have all the necessary materials. You will need: Concrete reinforcing wire mesh Look for mesh with a 6-inch by 6-inch grid and a gauge that feels robust. A common size is 150 feet long and 5 feet wide, which can make many cages. Wire cutters or bolt cutters: These are essential for cutting the thick wire mesh. Measuring tape: To ensure accurate cuts. Heavy-duty work gloves To protect your hands from sharp wire ends. Pliers (optional): For bending smaller wires.


Step 2: Determine Cage Dimensions and Cut the Mesh The size of your tomato cage depends on the variety of tomatoes you're growing. For most indeterminate varieties (which grow large), a cage that is 18-24 inches in diameter and 4-5 feet tall is ideal. To calculate the length of mesh needed for your desired diameter, multiply the diameter by π (~3.14). For an 18-inch diameter cage, you'll need roughly 56.5 inches of mesh. For a 24-inch diameter cage, you'll need about 75.4 inches. Unroll the wire mesh and, wearing your gloves, use the measuring tape to mark your desired length. With your wire cutters or bolt cutters, carefully cut the mesh. Make sure to cut just inside a vertical wire so that you have vertical wire ends sticking out along one edge – these will be used to secure the cage.

Step 3: Form the Cylinder Once you have your cut piece of mesh, it's time to form it into a cylinder. Carefully bring the two cut ends of the mesh together, overlapping them slightly. The vertical wire ends you left extended on one side will now align with the horizontal wires on the opposite end of the mesh.

Step 4: Secure the Cylinder With the ends overlapped, take each exposed vertical wire end and bend it around the corresponding horizontal wire on the opposite side of the mesh. Use your gloved hands or pliers to tightly secure each connection. Work your way around the entire cage, ensuring every connection is firm. This creates a strong, continuous cylinder. At this point, your basic cage structure is formed, but it's not yet stable.

Step 5: Create Anchor Points (Stakes) To keep the bottom of the cage firmly on the ground and prevent it from tipping over, especially with heavy fruit load, you'll need to add stake-like protrusions. Before fully closing the cage, identify several vertical wires near the base that will become your anchor points. Carefully uncoil the vertical wires at these spots, bending them outward and slightly downward. These extended wires will act as stakes when the cage is set in the ground.

Step 6: Finish Closing the Cage and Set it in Place Now, complete the enclosure of the cylinder by securing the remaining overlapped edges. Tighten all wire connections, including those for the anchor points. Ensure the cage is perfectly round and stable. Once your cage is fully assembled and sturdy, gently place it in the desired location in your garden. Push the anchor points firmly into the ground, making sure the cage stands upright and is stable.

 Can you teach me how to craft woven hay balls for rabbits, guinea pigs, and chinchillas?



Here's how you can craft woven hay balls for your small pets: First, gather your materials. You'll need some pet-safe hay (like timothy or orchard hay) and a bowl of water. You might also find a towel helpful for drying your hands.

Next, take a small handful of hay and moisten it slightly. You don't want it soaking wet, just pliable.

Now, form the damp hay into a small, loose ball. This will be the core of your woven hay ball.

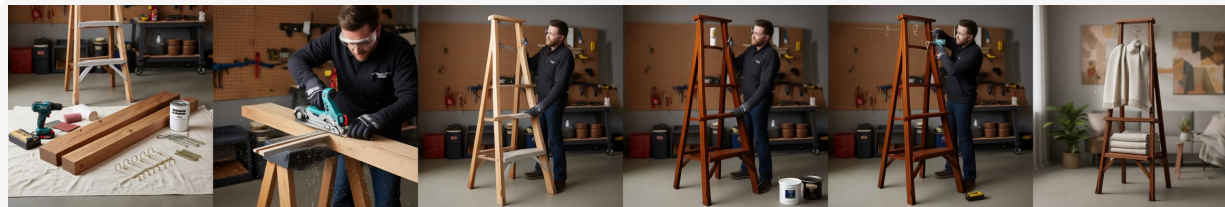
Then, take longer strands of hay, moisten them, and start weaving them around the core, tucking the ends in as you go. Imagine you're wrapping a gift, but with hay!

Continue adding layers and weaving the hay until you have a firm, compact ball. The water helps the hay stick together, allowing for a tight weave.

Finally, place your finished hay balls in a warm, dry place to air dry completely. This can take a day or two, depending on the humidity. Once fully dry, they'll be firm and ready for your pets to enjoy!

Figure A3. Generation results of DuoGen on How-to-500.

 Can you teach me how to turn a wooden ladder into a blanket rack?



Here's how you can transform a wooden ladder into a stylish and functional blanket rack! First, gather your materials. You'll need a wooden ladder (preferably with a nice rung design), sandpaper, wood stain or paint, clear coat sealant, and hooks. A drill is also necessary for installing the hooks.


Start by de-ringing the ladder. This means carefully removing the wooden rungs that connect the side rails. You can use a reciprocating saw or a jigsaw to cut through the nails or screws holding the rungs in place. Be sure to wear safety glasses!

Once the rungs are removed, you'll have an open ladder structure. Now it's time for preparation. Use sandpaper to lightly scuff up the entire ladder. This helps the stain or paint adhere better. If you want to restore the natural wood look, use a finer-grit sandpaper.

Next, apply your chosen finish. If you're using wood stain, apply it evenly with a brush or rag, allowing it to dry according to the manufacturer's instructions. For paint, apply thin, even coats, allowing each coat to dry before applying the next. A clear coat sealant can be applied after the stain or paint has dried completely to provide extra protection and a nice sheen.

Finally, attach the hooks. Measure and mark where you want your hooks to go. Typically, spacing them about 6-10 inches apart works well for hanging blankets. Use a drill to create pilot holes and then screw in the hooks securely.

Your beautiful new blanket rack is now complete and ready to hold your favorite throw blankets!

 Can you teach me how to feed bees sugar syrup in early spring?



Feeding bees sugar syrup in early spring can be a crucial step to help them build up their colony after winter and prepare for the upcoming nectar flow. Here's a step-by-step guide: Step 1: Assess the Colony's Needs Before you start feeding, it's important to check your hive. Look for signs of low honey stores and a relatively small population. If the bees have plenty of honey left from last year, they might not need feeding. Early spring feeding is primarily for stimulation and supplementing stores.


Step 2: Prepare the Sugar Syrup For early spring feeding, a 1:1 sugar-to-water ratio is generally recommended. This ratio stimulates the queen to lay more eggs, helping to grow the colony. Use plain white granulated sugar. Avoid brown sugar or other types as they can be harmful to bees. Heat the water, then dissolve the sugar completely. Let it cool before feeding.

Step 3: Choose a Feeder Type There are several types of feeders you can use: Entrance feeders: These are external and easy to refill, but can attract robbing bees from other colonies. Boardman feeders: Similar to entrance feeders, they attach to the hive entrance. Internal frame feeders: These sit inside the hive like a frame and are less prone to robbing. Top feeders (or hive-top feeders): These sit above the inner cover and are often the easiest and safest for both bees and beekeeper.

Step 4: Install the Feeder For a top feeder, place it directly on top of the inner cover, ensuring the bees have access to the syrup. If using an entrance or internal frame feeder, install it according to the manufacturer's instructions. Ensure there are no leaks and that the bees can safely access the syrup without drowning; floats (like twigs or straw) can be placed in open feeders for this purpose.

Step 5: Fill the Feeder with Syrup Carefully pour the cooled sugar syrup into the chosen feeder. Avoid spilling syrup outside the hive, as this can attract ants or other pests, and encourage robbing by other bee colonies.

Step 6: Monitor and Refill Check the feeder regularly (every few days) and refill as needed. Bees can consume a surprising amount of syrup quickly. Continue feeding until the natural nectar flow begins, or when you notice the bees are no longer taking the syrup. Stop feeding once the main nectar flow starts to ensure your honey crop is not diluted with sugar syrup.

 Can you teach me how to dye pasta in rainbow colors for crafts?



Dyeing pasta in rainbow colors is a fun and easy way to create vibrant materials for various crafts! Here's how you can do it: Step 1: Gather Your Materials You'll need dry pasta (any shape will work, but larger, more intricate shapes tend to show the color better), rubbing alcohol or white vinegar, food coloring, shallow trays or bowls, and paper towels for drying.

Step 2: Prepare Your Work Area Lay out some paper towels or a drop cloth to protect your surface from the dye and pasta. Have your trays of rubbing alcohol or white vinegar ready.

Step 3: Add Food Coloring to the Dyeing Liquids Pour a small amount of rubbing alcohol or white vinegar into each tray. Then, add several drops of food coloring to each tray, choosing different colors to create your rainbow. Stir gently until the liquid is evenly colored.

Step 4: Dye the Pasta Add a handful of pasta to one of the trays with the chosen color. Toss it gently to ensure the pasta pieces are fully coated. Lift them out, letting any excess dye drip back into the tray.

Step 5: Repeat for Different Colors Move the dyed pasta to a clean paper towel to dry slightly, then transfer it to a tray with a different colored liquid. Repeat this process for all the colors of the rainbow until your pasta is dyed in beautiful rainbow hues.

Step 6: Dry and Use Your Dyed Pasta Spread the rainbow-colored pasta on a baking sheet lined with paper towels and let it dry completely. Once fully dry, your pasta is ready for crafting! It will be perfect for making jewelry, bracelets, or any other creative project where you want to add a splash of colorful pasta.

Figure A4. Generation results of DuoGen on How-to-500.

 Can you teach me how to make handmade paper with embedded seeds for plantable cards?



That's a fantastic idea! Making handmade paper with embedded seeds for plantable cards is a creative way to give back to the earth. Here's how you can do it: **Step 1: Gather Your Materials** You'll need a deconstructing kit (old books, papers), a blender, deckle and screen (you can make these yourself with fine-meshed screens and picture frames), seed packets, water, and a flat tray or tub to work in.

Step 2: Deconstruct Your Source Material Carefully take apart old books or papers. This can be the most tedious part but also the most rewarding. You want to tear the pages into manageable strips, trying to avoid large chunks of text if you prefer a more "blank" sheet.

Step 3: Wet Your Material Place your torn paper strips into a bowl of water. Agitate them gently until the fibers begin to separate and the paper becomes mushy. This usually takes about 10-15 minutes in a blender, but can be done by hand as well with vigorous stirring.

Step 4: Add the Seeds Once your fibers are sufficiently pulpy, scoop some of the mixture into another container and add a small amount of seeds. A pinch per handful of paper pulp should be enough. Mix gently to distribute the seeds evenly.

Step 5: Form Your Paper Sheets Pour a shallow layer of your seed-infused pulp onto your deckle (the inner frame works with the screen). Remove excess water by scraping the edges with a rubber squeegee, then lift off the deckle, letting the excess water drain through the screen.

Step 6: Press and Dry Your Paper Place your screen on top of a flat surface and gently press down with a roller or another frame to remove more water. You can also use blotter paper or even towels to absorb the remaining moisture. Once sufficiently dry, carefully peel the paper off the screen.

 Can you teach me how to build a multi-station bird feeding post for your backyard?



Here's how you can build a multi-station bird feeding post, providing a delightful spot for your feathered friends in your backyard! First, gather all your materials. You'll need a sturdy post (like a 4x4 wooden post, about 5-6 feet tall), several smaller perching posts or feeders (you can buy these or fashion them from branches), screws or nails, a drill, a saw, and possibly some birdseed.

Next, prepare your main post. If it's not already at the desired height, use a saw to cut it to size. Ensure one end is perfectly square. Then, dig a hole that's slightly wider and deeper than your post.

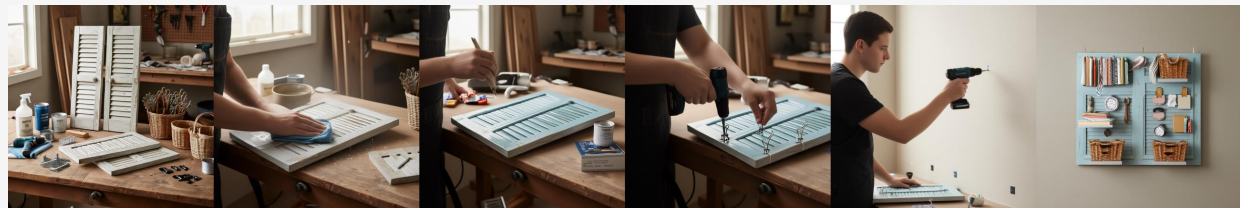
To stabilize the post, line the hole with concrete or gravel. This will prevent the post from wobbling once it's in place. Place the post into the hole, making sure it's straight, and fill the remaining space with more concrete or gravel around the base.

Allow the concrete to dry completely. Once set, you can add small platforms or hooks to the main post where you can attach your various feeder stations. These platforms can be simple pieces of wood cut to size.

Now it's time to attach your feeder stations! You can create different types of feeders. Some can be simple branches with drilled holes to hold seeds, others could be commercial suet feeders, or Thistle Sock feeders hung with hooks. Attach these to your platforms on the main post.

Finally, fill your new feeders with birdseed and watch as your backyard visitors discover their new feeding station! Regularly clean and refill the feeders to keep the area inviting for birds.

 Can you teach me how to repurpose wooden shutters into wall organizers?



Repurposing wooden shutters into wall organizers is a fantastic way to add rustic charm and practical storage to your home. Here's how you can do it: **Step 1: Gather Your Materials** First, you'll need your wooden shutters. You can find these at flea markets, antique shops, or even salvage yards. Look for shutters that are sturdy and in a style you like. You'll also need cleaning supplies, sandpaper, paint or stain (optional), a drill, screws, picture hangers or D-rings, and various hooks, small baskets, or clips for organizing.

Step 2: Clean and Prepare the Shutters Thoroughly clean your shutters to remove any dust, dirt, or debris. You can use a mild cleaner and a cloth. If the shutters are very dirty or have wear and tear, sand them down to smooth out rough spots and remove old finishes. This prepares the surface for painting or staining.


Step 3: Paint or Stain (Optional) If you want to change the color of your shutters, now's the time to paint or stain them. Choose a color that complements your home decor. Apply an even coat and let it dry completely according to the product instructions. You might need two coats for full coverage.

Step 4: Attach Hanging Hardware Once the paint or stain is dry, attach picture hangers or D-rings to the back of the shutters. Make sure they are evenly spaced and positioned high enough to hold the weight of the shutter when closed. If the shutters are heavy, consider using two sets of hangers for better support.

Step 5: Drill Mounting Screws into the Wall Now, carefully mark the desired hanging position on your wall. It's best to have someone help you hold the shutters up against the wall while you mark the corresponding spots on the wall. Once marked, drill pilot holes for your mounting screws. Use screws that are long enough to securely hold the hangers and shutters, but not so long that they go through the face of the shutter.

Step 6: Add Organizing Elements Finally, adorn your newly mounted shutter organizer! Use various hooks to hang small bags, scarves, hats, or even utensils. Small baskets can hold books, notebooks, or other small items. Clips can be used to display photos or art. Arrange them creatively to suit your needs and aesthetic.

Figure A5. Generation results of DuoGen on How-to-500.

 Can you teach me how to clean and deodorize indoor trash cans?



Cleaning and deodorizing indoor trash cans is a simple process that can significantly improve the hygiene and freshness of your home. Here's a step-by-step guide: **Step 1: Empty the Trash Can First**, ensure the trash can is completely empty. Remove the trash bag and dispose of any remaining debris inside the can.

Step 2: Rinse the Can Take the trash can outdoors or into a large utility sink. Rinse it thoroughly with water to remove any loose dirt, food particles, or residue.

Step 3: Apply Cleaning Solution Create a cleaning solution. You can use a mixture of warm water and dish soap, or for a more potent clean, add a cup of white vinegar or a few tablespoons of baking soda to the water. Scrub the inside and outside of the can with a brush or sponge, paying attention to any stubborn stains or grime.

Step 4: Rinse Thoroughly Again Once you've scrubbed the entire can, rinse it thoroughly with clean water to remove all traces of the cleaning solution. Make sure no soapy residue is left behind.

Step 5: Dry Completely It's crucial to dry the trash can completely to prevent mold and mildew growth, which can cause odors. You can air-dry it in the sun, wipe it down with a clean cloth, or turn it upside down to allow water to drain.

Step 6: Deodorize To deodorize, place a box of baking soda, a few coffee grounds, or a silica gel packet at the bottom of the empty trash can. These absorb odors effectively. When full, replace the deodorizer or refresh it as needed.

 Can you teach me how to build a spa gift basket with homemade bath salts for Mother's Day?



Here's a guide to creating a lovely spa gift basket with homemade bath salts for Mother's Day! First, let's gather our supplies for the bath salts. You'll need Epsom salt, coarse sea salt, baking soda, essential oils (like lavender or eucalyptus), and optional dried herbs or flower petals for extra visual appeal.

Next, in a large bowl, combine 2 cups of Epsom salt, 1 cup of coarse sea salt, and 1/2 cup of baking soda. The baking soda helps to soften the water and cleanse the skin.

Now, add about 20-30 drops of your chosen essential oil to the salt mixture. Stir everything together thoroughly to ensure the essential oil is evenly distributed. If using, mix in a handful of dried herbs or petals.

Once your bath salts are mixed, transfer them into an airtight jar. This will keep them fresh and preserve their fragrance until Mother's Day.

Now, for the gift basket! Choose a nice basket and fill it with crinkled paper or tissue paper for cushioning. Place the homemade bath salts prominently in the basket.

Finally, add other spa-themed items to complete the basket. Consider a small loofah, a bath sponge, a few candles, a bottle of lotion, and perhaps a small brush for exfoliation. Arrange everything appealingly for a beautiful and relaxing Mother's Day gift!

 Can you teach me how to create foil stickers with a Minc machine?



Here's how to create foil stickers using a Minc machine, step-by-step: **Step 1: Design Your Sticker** First, you'll need to design your sticker. This can be done using graphic design software like Adobe Illustrator, CorelDRAW, or even online tools like Canva. Make sure your design is vector-based for the best results. Consider the size and shape of your final sticker.

Step 2: Cut Your Sticker Image from Adhesive Vinyl Once your design is ready, load adhesive vinyl into your cutting machine (e.g., a Silhouette Cameo, Cricut, or Minc). Use your design file to cut out the exact shape of your sticker. For a clean cut, make sure the outer edges of your design match the desired final sticker size.

Step 3: Prepare Your Foil Unroll your heat transfer foil. There are two main types: shiny and matte. Shiny foil will give your sticker a vibrant metallic finish, while matte foil will cover the vinyl entirely without shine. For this guide, we'll assume matte foil. Position the foil so it covers the entire cut sticker image.

Step 4: Load into the Minc Machine Carefully place both the cut vinyl sticker and the foil onto the Minc machine's rollers. Ensure the vinyl is right-side up, facing the first roller. The Minc machine uses heat and pressure to transfer the foil onto the vinyl.

Step 5: Run Through the Minc Machine Press the "GO" button on your Minc machine. It will automatically apply the correct amount of heat and pressure to transfer the foil. You'll see the foil gradually change color as it bonds to the vinyl.

Step 6: Peel and Enjoy! Once the Minc machine finishes, carefully peel off the excess foil and unadhered vinyl. You should be left with your beautiful, custom foil sticker!

Figure A6. Generation results of DuoGen on How-to-500.

How to Cook Eggs*:
Eggs are packed with protein and are a truly versatile food that can be cooked in many ways. Here are just a few of the most popular methods.
You need to assist human user to complete this task via scrambled. Given the previous steps, you need to predict the subsequent 3 steps to help the user to finish the task. The previous steps are:
Spray a medium skillet with cooking spray.



Whisk together the eggs, milk, salt, and pepper.



Here are the next three steps to cook your eggs:

Add the whisked egg mixture to the skillet.



Gently stir the eggs with a spatula until they are cooked to your desired consistency.



Remove the skillet from the heat and serve immediately.



How to Drink Milk if You Hate It*:
You need to assist human user to complete this task via mixing milk with other foods. Given the previous steps, you need to predict the subsequent 4 steps to help the user to finish the task. The previous steps are:
Dunk bread or cookies into milk.

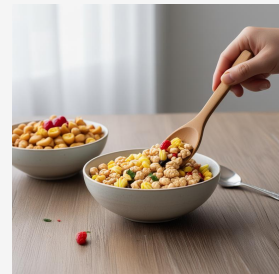


Sure, here are four steps to help someone who dislikes milk incorporate it into their diet by mixing it with other foods:

Step 1: Create Milkshake Mixtures
Add milk to smoothies, milkshakes, or frosties to enhance flavor and increase dairy intake.



Step 2: Incorporate into Cereal or Oatmeal
Mix milk into cereals or oatmeal to subtly integrate it into your morning meal.



Step 3: Use in Baked Goods
Incorporate milk into baked goods like pancakes, waffles, or muffins for a hidden way to consume dairy.



Step 4: Make Sauces and gravies
Blend milk into sauces and gravies to enrich their flavor and texture, making them less noticeable.

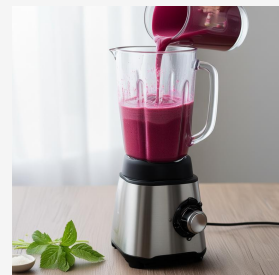


Figure A7. Generation results of DuoGen on InterleavedBench.

How to Clean Wooden Bowls: Wooden bowls add warmth and style to your decor or dining room table settings, but they require different care than standard glass or plastic dishes. With regular washing and occasional treatments with oil, wooden bowls will last years. When you have tough stains or buildup, try salt and lemon juice or buff your bowls gently with sandpaper.



You need to assist human user to complete this task via treating stains and buildup. Given the previous steps, you need to predict the subsequent 2 steps to help the user to finish the task. The previous steps are:

Give wooden bowls a deep clean with salt and lemon.



To effectively clean your wooden bowl, follow these steps:

****Step 1: Rinse the Bowl****

First, rinse the wooden bowl thoroughly under warm running water to remove any loose debris, food particles, or surface dirt. This initial rinse helps prepare the bowl for deeper cleaning.



****Step 2: Apply Salt and Lemon Juice****

After rinsing, sprinkle coarse salt evenly over the surface of the wet bowl. Then, squeeze fresh lemon juice directly onto the bowl, ensuring the entire surface is coated. The salt acts as an abrasive to scrub away stains, while the lemon juice provides natural disinfection and helps break down grime. Let it sit for about 10-15 minutes.



****Step 3: Scrub and Rinse Thoroughly****

Using a soft-bristled brush or a clean cloth, gently scrub the bowl in circular motions to work the salt and lemon juice into the wood grain. Pay extra attention to stubborn areas. Once you've scrubbed all surfaces, rinse the bowl completely again with warm water until all the salt and lemon residue is gone. Make sure there are no lingering grains of salt, as they can damage the wood during drying.



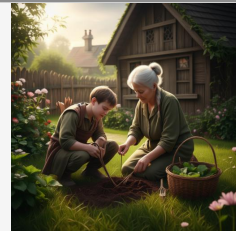
Generate a story based on the following context on boy mitcham milly .



In the quaint village of Milly, there's an legendary story surrounding a special type of mint called "Mintcham Milly." This mint is said to possess magical properties, making any beverage it's added to incredibly refreshing. Once upon a time, in a small garden at the edge of Milly, a wise old woman named Elise tended to a patch of this sacred mint. Every summer, she carefully harvested the mint, ensuring its purity and strength.



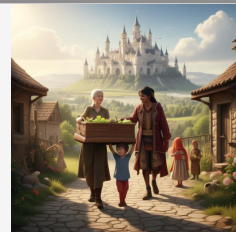
Elise shared her knowledge with young Boy Mitcham, who became her apprentice. He learned how to identify the perfect conditions for growing the mint, from fertile soil to the right amount of sunlight.



As Mitcham grew older, he took over the care of the mint patch. Word spread throughout the village about the extraordinary cooling effects of his grandmother's Mintcham Milly. People would gather around as Mitcham brewed a refreshing cup for them to taste.



One hot summer, a traveling merchant stopped in Milly and was captivated by the village's vitality. Impressed by Mitcham's skill, the merchant offered him a generous sum to bring some of the mint back with him to his far-flung empire. Mitcham, true to his word, packed the precious leaves into a wooden box and set off on his journey.



Upon arriving in the merchant's grand city, Mitcham sold the Mintcham Milly at an exorbitant price, believing it was the rarest and most valuable commodity. However, the real treasure lay not in the profit, but in the joy and refreshment it brought to countless thirst-quenching moments. And so, the legend of Mintcham Milly continued to thrive, connecting people across distant lands.

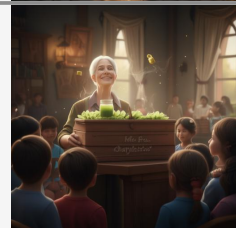


Figure A8. Generation results of DuoGen on InterleavedBench.

2. Image Generation and Editing Examples

Fig. A9, Fig. A10, and Fig. A11 shows more examples on image generation, image editing, and multi-reference image generation, respectively.

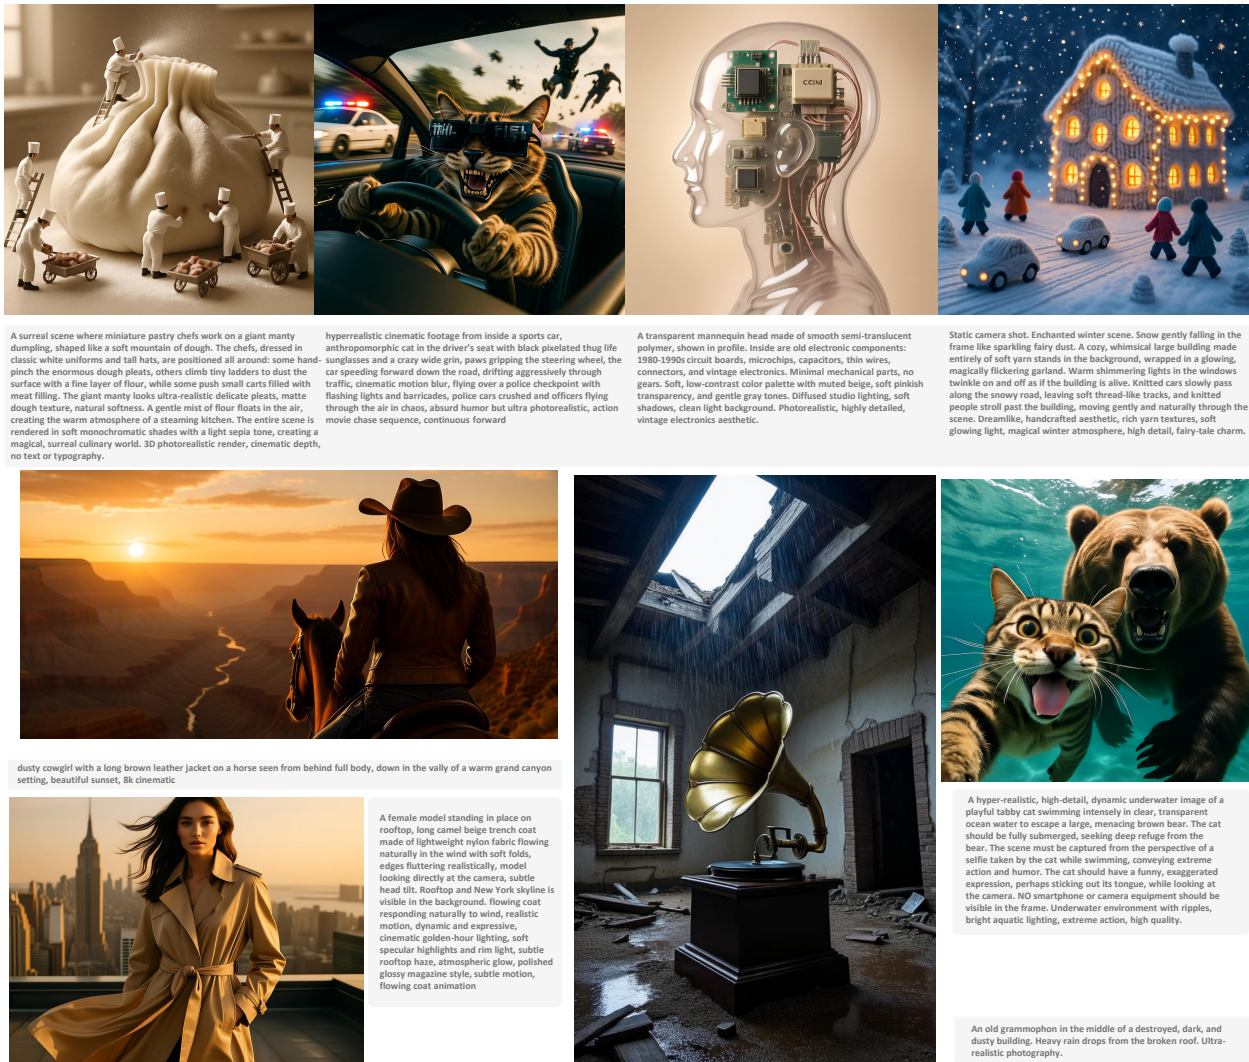


Figure A9. Text-to-Image generation results of DuoGen.

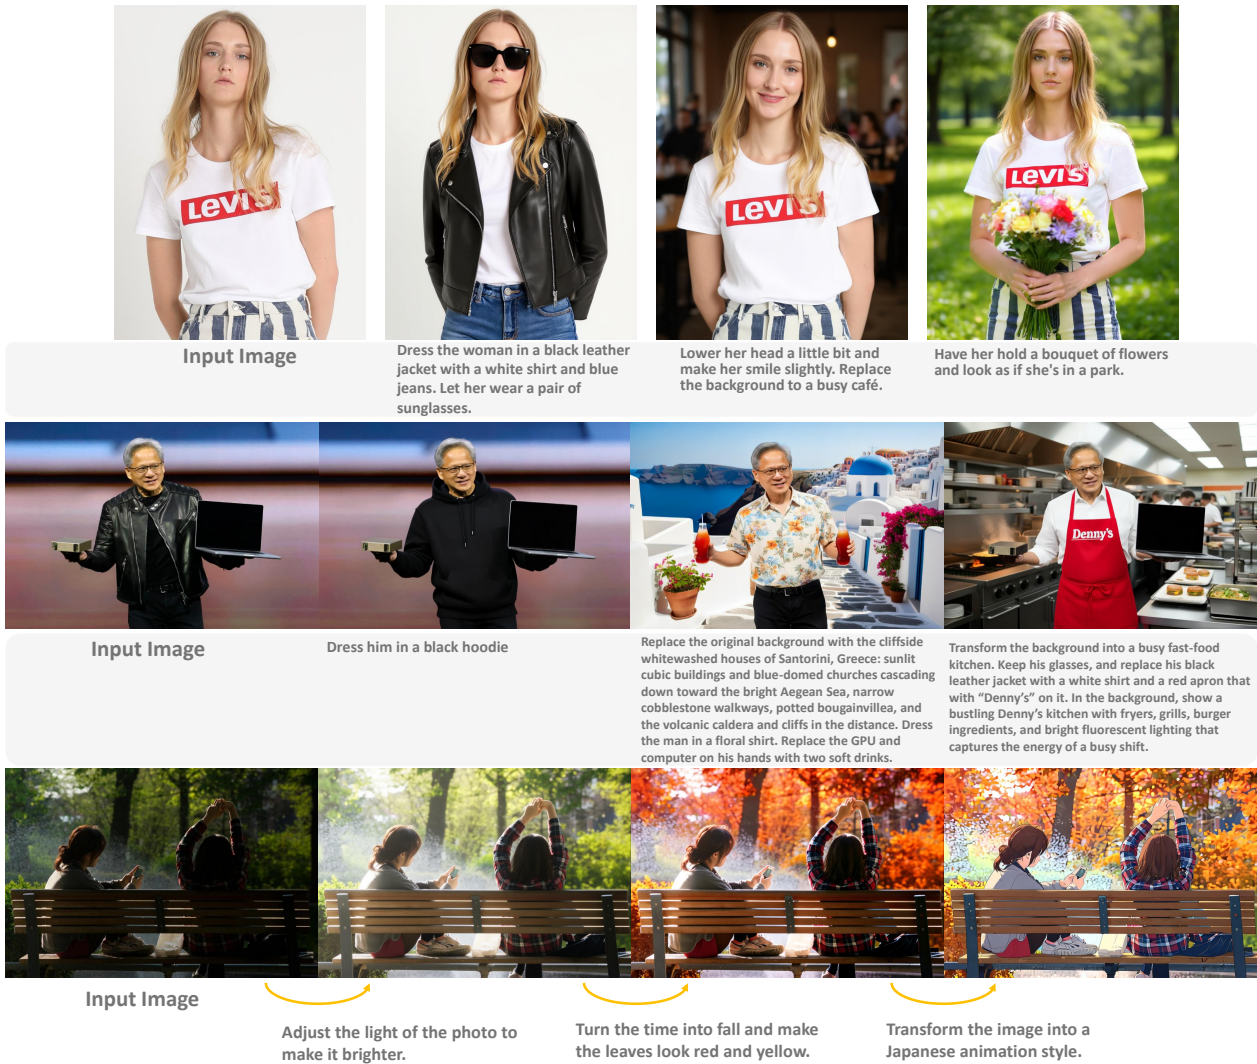


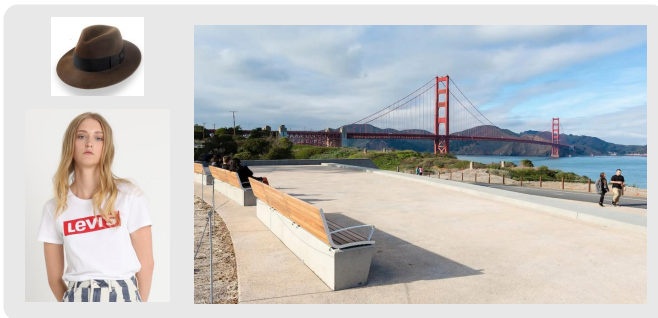
Figure A10. Image editing results of DuoGen.



Put the woman in a tense situation inside what appears to be a convenience store. Her stance is firm and focused, looking directly at the camera as she holds a gun pointed directly at the camera, suggesting a confrontation or a high-stakes scenario. The setting is a typical convenience store with shelves stocked with various products, including snacks and beverages. The overall atmosphere is charged with tension, hinting at a dramatic or action-packed moment captured in the scene.



Let the woman wear this white jacket.



Let the woman wear the fedora in the second image. Replace the background with the third image.



Let the woman wear the glasses and using the laptop to study at the library.

Figure A11. Multi-reference image generation results of DuoGen.

3. Comparison and Results of More VLM Judge

Table A1. **Comprehensive Comparison of Automated MLLM Judges versus Human ELO Ratings on Interleaved Generation Tasks.** T-Com, I-Com, I-Co, I-Q, and IT-Co denote text completeness, image completeness, image-coherence, image quality, and image-text coherence, respectively. This table expands on main results by introducing GPT-5 and GPT-4o as additional judge models. Qwen and Human ELO scores are directly pulled from our human study.

Model	T-Com				I-Com				I-Co				I-Q				IT-Co			
	GPT5	GPT4o	Qwen	ELO	GPT5	GPT4o	Qwen	ELO	GPT5	GPT4o	Qwen	ELO	GPT5	GPT4o	Qwen	ELO	GPT5	GPT4o	Qwen	ELO
Nano Banana [11]	4.54	4.84	4.83	1393	4.01	4.89	4.60	1369	4.32	4.97	4.56	1420	4.65	4.99	4.94	1473	4.53	4.99	4.81	1419
SEED-LLaMA [9]	1.82	3.37	2.39	940	1.62	3.43	1.68	963	2.22	4.34	2.13	996	3.16	3.58	3.72	961	2.31	3.61	1.76	942
MiniGPT-5 [40]	1.87	2.54	2.17	1161	1.88	2.74	2.06	1111	1.99	2.43	1.82	1058	2.72	2.81	3.68	1071	2.43	2.35	2.22	1135
Zebra-CoT [16]	2.89	3.00	3.43	1078	2.34	3.57	3.70	1115	2.80	4.28	3.73	1096	3.57	4.31	4.42	1111	3.26	4.30	4.04	1256
DuetGen	4.18	4.56	4.79	1428	4.05	4.85	4.64	1442	4.07	4.87	4.54	1430	4.50	4.96	4.92	1384	4.42	4.95	4.78	1402

Table A2. **Agreement and Correlation between MLLM Judges and Human Preference.** “Agr.” denotes the agreement rate, and “Corr” represents the Pearson correlation coefficient.

MLLM Judge	T-Com		I-Com		I-Co		I-Q		IT-Co	
	Agr.	Corr	Agr.	Corr	Agr.	Corr	Agr.	Corr	Agr.	Corr
GPT-5	92.9%	0.87	97.2%	0.95	91.0%	0.89	90.7%	0.95	95.9%	0.89
GPT-4o	91.2%	0.85	98.5%	0.91	87.8%	0.85	92.2%	0.96	92.5%	0.83
Qwen3VL	94.3%	0.82	98.3%	0.89	86.2%	0.83	91.8%	0.90	90.0%	0.78

We further supplement our analysis by incorporating additional MLLMs as judges. Specifically, we employ GPT-4o and Qwen3-VL-32B-Instruct under identical evaluation prompts. The comparative results are summarized in Table A1. Across the primary metrics (T-Com, I-Com, I-Co, I-Q, IT-Co), we observe strong cross-judge consistency in rank ordering. Notably, Nano Banana and our proposed DuoGen are consistently identified as top-performing methods across nearly all metrics, irrespective of whether evaluation is conducted by human annotators or automated MLLM judges.

Table A2 presents correlations with human judge. Open-source models such as Qwen exhibit strong alignment with human and proprietary judges on fundamental metrics (e.g., T-Com and I-Com). In contrast, GPT-5 demonstrates greater sensitivity to subtle quality variations, particularly in nuanced visual reasoning metrics such as global visual coherence (I-Co and IT-Co).

4. Detailed Results on Image Generation and Editing Benchmarks

Due to space limitations in the main paper, we reported only the overall scores for the GenEval [10] and ImgEdit [38] benchmarks. Here, we provide the detailed scores for each subtask. Table A3 presents the complete GenEval results, and Table A4 provides the detailed results for ImgEdit.

Model Type	Method	Single Object	Two Object	Counting	Colors	Position	Attribute Binding	Overall
Commercial	GPT-4o-Image [21]	0.99	0.92	0.85	0.92	0.75	0.61	0.84
Generation	SDXL [24]	0.98	0.74	0.39	0.85	0.15	0.23	0.55
	DALLE-3 [20]	0.96	0.87	0.47	0.83	0.43	0.45	0.67
	FLUX.1-dev [14]	0.98	0.93	0.75	0.93	0.68	0.65	0.82
	Qwen-Image [31]	0.99	0.92	0.89	0.88	0.76	0.77	0.87
Unified Model	Emu3 [29]	0.98	0.71	0.34	0.81	0.17	0.21	0.54
	Show-o [34]	0.95	0.52	0.49	0.82	0.11	0.28	0.53
	Janus-Pro-7B [4]	0.99	0.89	0.59	0.9	0.79	0.66	0.8
	MMaDA [36]	0.99	0.76	0.61	0.84	0.2	0.37	0.63
	MetaQuery-XL* [23]	-	-	-	-	-	-	0.8
	BLIP3o [3]	-	-	-	-	-	-	0.84
	Bagel [6]	0.99	0.94	0.81	0.88	0.64	0.63	0.82
	UniWorld-V1 [17]	0.99	0.93	0.79	0.89	0.49	0.7	0.8
OmniGen2 [32]	1	0.95	0.64	0.88	0.55	0.76	0.8	
Unified Interleaved Generation Model	Uni-CoT** [26]	0.99	0.96	0.84	0.92	0.57	0.71	0.83
	Ours	0.82	0.99	0.94	0.91	0.84	0.80	0.88

Table A3. **Detailed comparison on GenEval benchmark.** * denotes that the model is using LLM to rewrite prompts. ** Uni-CoT is using interleaved generation to improve image generation quality.

Model Type		Add	Adjust	Extract	Replace	Remove	Background	Style	Hybrid	Action	Overall
Commercial Model	GPT-4o-Image [21]	4.61	4.33	2.90	4.35	3.66	4.57	4.93	3.96	4.89	4.20
Generation	ICEdit [39]	3.58	3.39	1.73	3.15	2.93	3.08	3.84	2.04	3.68	3.05
	Step1X-Edit [19]	3.88	3.14	1.76	3.40	2.41	3.16	4.63	2.64	2.52	3.06
	FLUX.1 Kontext [Pro] [15]	4.25	4.15	2.35	4.56	3.57	4.26	4.57	3.68	4.63	4.00
	Qwen-Image [31]	4.38	4.16	3.43	4.66	4.14	4.38	4.81	3.82	4.69	4.27
Unified Model	OmniGen [33]	3.47	3.04	1.71	2.94	2.43	3.21	4.19	2.24	3.38	2.96
	Bagel [6]	3.56	3.31	1.70	3.30	2.62	3.24	4.49	2.38	4.17	3.20
	UniWord-V1 [17]	3.82	3.64	2.27	3.47	3.24	2.99	4.21	2.96	2.74	3.26
	OmniGen2 [32]	3.57	3.06	1.77	3.74	3.20	3.57	4.81	2.52	4.68	3.44
	OVIS-U1 [28]	4.13	3.62	2.98	4.45	4.06	4.22	4.69	3.45	4.61	4.00
	Lumina-DiMOO [35]	3.82	-	-	3.83	2.76	-	4.18	-	-	-
Interleaved	DuoGen	4.53	4.33	2.28	4.69	4.71	4.61	4.51	3.85	4.67	4.19

Table A4. **Detailed comparison on ImgEdit benchmark.**

5. Details of the Interleaved Generation Benchmark

Below is the prompt used to evaluate model outputs on our interleaved generation benchmark.

Evaluation Prompt for Our Interleaved Generation Benchmark (GPT 5)

You are an expert evaluator specializing in assessing the generation quality of multi-modal interleaved outputs (i.e., sequences of text and images).
Your task is to evaluate a model's generated response to a user request that may include both text and image inputs.

****Input Format****

- ****User Input:**** A user question or request, possibly with reference images.
- ****Model Output:**** The model's response consisting of interleaved text and generated images illustrating steps, processes, or results.

You must evaluate the model output across the following dimensions, assigning scores ****from 1 to 5**** according to the detailed criteria.

Whole-Sequence Metrics

****1. Text Completeness****

Assess whether the generated text fully and coherently addresses the user's request.
A complete response includes all necessary reasoning steps and provides a clear final result.

Score	Description
5	Fully complete; all necessary steps and reasoning are included with a clear final answer.
4	Mostly complete; all major points are covered, though some minor steps may be missing.
3	Reaches a final answer but omits one or more **critical steps** or contains logical gaps.
2	Only partially complete; major reasoning steps are missing or incoherent.
1	Irrelevant or fails to address the user's request.

****2. Image Completeness****

Evaluate whether the images collectively cover all important parts of the answer.
Ideally, each major step or concept in the text should be visually illustrated.
(For simple questions with few steps, this criterion may be relaxed.)

Score	Description
5	Every major step or part is illustrated with an appropriate image.
4	Most major steps are illustrated; one or two non-critical parts are missing images.
3	One or two **major** steps lack visual representation.
2	Most steps lack images; only a few parts are visually supported.
1	Fewer than two relevant images are included.

****3. Image Coherence****

Evaluate visual consistency across all generated images.
Objects, characters, and scenes should remain coherent when continuity is expected, unless explicitly changed by the text. Do not remove objects should be remained in the previous images.

Score	Description
5	Excellent coherence; consistent scenes and objects across all steps.
4	Generally coherent with minor inconsistencies (e.g., small detail changes).
3	Mostly coherent but with one or two **notable inconsistencies** (e.g., major object or scene changes).
2	Low coherence; multiple inconsistencies in key objects or settings.
1	Very poor coherence; scenes change arbitrarily without justification.

Per-Image Metrics

Evaluate each generated image individually along the following two dimensions.

1. Image Quality

Assess the visual realism, clarity, and overall composition of each image.

Look for common generative flaws such as distorted hands, incorrect text, or visual artifacts.

Score	Description
*****	Highly realistic and visually coherent; no major artifacts.
4	Natural and clear; only minor imperfections or artifacts.
3	Generally clear, but noticeable artifacts or rendering flaws exist.
2	Major artifacts or blurriness; multiple unclear elements.
1	Heavily distorted or unrecognizable image.

2. Image-Text Coherence

Evaluate how well each image visually represents its accompanying text segment.

The image should clearly depict the described object, action, or result.

Score	Description
*****	Perfect alignment; all key objects, actions, and details match the text.
4	Strong alignment with only minor omissions or mismatches.
3	Generally matches, but misses multiple described elements.
2	Loosely related; only some overlap with the text.
1	Irrelevant or repeated image; no clear relation to text.

Output Format

Please provide your evaluation in **strict JSON** format as follows:

```
{
  "scores": {
    # whole-sequence metrics
    "text-completeness": {
      "Justification": "2-3 sentence explanation of strengths or weaknesses.",
      "Score": int (1-5)
    },
    "image-completeness": {
      "Justification": "2-3 sentence explanation of strengths or weaknesses.",
      "Score": int (1-5)
    },
    "image-coherence": {
      "Justification": "2-3 sentence explanation of strengths or weaknesses.",
      "Score": int (1-5)
    },
    # per-image metrics
    "image-quality": {
      "image_0": {
        "Justification": "2-3 sentence explanation of visual quality.",
        "Score": int (1-5)
      },
      "image_1": {
        "Justification": "...",
        "Score": int (1-5)
      },
      ... # list all the images in the response
    },
    "image-text-coherence": {
      "image_0": {
        "Justification": "2-3 sentence explanation of image-text match.",
        "Score": int (1-5)
      },
      "image_1": {
        "Justification": "...",
        "Score": int (1-5)
      },
      ... # list all the images in the response
    },
  }
}
```

Now evaluate the following conversation:

User Request: *{user_request}*

Model Response: *{model_response}*

6. Data Engine for Webpages

As discussed in Sec. 3.1 of the main paper, we convert raw web pages into multi-modal conversations using a data engine with a series of cleaning and rewriting steps. The full pipeline consists of four major stages: 1) rewriting and splitting the raw interleaved content; 2) captioning and categorizing all images; 3) removing duplicate images and reordering the interleaved text–image sequence; and 4) converting the cleaned interleaved content into a user–assistant conversation.

Below, we present the prompt used in the first stage, which transforms the raw webpage into introduction and main content, both represented as interleaved image–text sequences. In this step, the VLM also rewrites the content to correct errors and remove irrelevant elements such as advertisements. To improve throughput, we replace image content with text placeholders, as this step focuses solely on text rewriting rather than image understanding.

Prompt for Cleaning Raw Webpages (Qwen2.5 VL 72B)

You are a helpful data annotator. Your task is to clean and refine noisy interleaved image-text sequences collected from websites.

Input format:

```
[
  "text1",
  "image1",
  "text2",
  "image2",
  ...
]
...
```

Instructions:

1. Clean text:
 - Remove irrelevant strings, HTML tags, special characters, and extra whitespace.
2. Ensure coherence:
 - Revise text so it is clear, logical, and consistent with the associated image.
 - Reconstruct unclear text if necessary.
3. Remove noise:
 - Delete advertisements, navigation links, or unrelated content.
 - The final sequence must be self-contained.
4. Empty output rule:
 - If the entire sequence is an advertisement or lacks meaningful content, return an empty list: []
5. Split introduction and content:
 - If the input list can be divided into an introduction part introducing the background of final results for the rest of the part, split the input list into an "introduction" part and a "content" part.

Note: the image is just a placeholder, you cannot see the image. You cannot modify the position of the image or delete the image from the sequence.

Output a json object with the following format:

```
{
  "introduction": [
    "refined text1",
    "image1",
    "refined text2",
    "image2",
    ...
  ],
  "content": [
    "refined text3",
    "image3",
    "refined text4",
    "image4",
    ...
  ]
}
```

Now the input is:

```
[
  {input_sequence}
]
```

Below is the prompt used to caption and classify the images in the interleaved sequence.

Prompt for Image Caption and Classification in Interleaved Content (Qwen2.5 VL 72B)

Generate an image description for the given image with the following requirements:

- Descriptions:
 - A short description (20–30 words).
 - A detailed description (around 50 words).
- Image type classification: Choose one of the following categories:
 - photo realistic
 - comic
 - 3D model
 - UI (for mobile app or computer software interfaces)
 - others
 - Invalid image that are not not meaningful: meaning less icons, advertisements with links, QRcode, etc.
- Description rules:
 - Always begin with "A image of ..." or "A image showing ...".
 - Do not start with "This image shows ...".

- Output format:

```
{
  "detailed_description": "...",
  "short_description": "...",
  "image_type": "..."
}
```

Input Image

<Base64 string of the image>

Below is the prompt used to remove duplicate images and to reorder the interleaved image and text chunks.

Prompt for Reordering and Removing Images in the Interleaved Sequence (Qwen2.5 VL 72B)

You are a helpful data annotator. Your task is to reorder the place of images and text chunks in an interleaved image-text sequences collected from websites. The input will be a two interleaved sequence from the introduction part and the content part.

Introduction part:

```
[
  "text1",
  "image1",
  "text2",
  "image2",
  ...
]
```

Content part:

```
[
  "text3",
  "image3",
  "text4",
  "image4",
  ...
]
```

Instructions:

1. Remove irrelevant images:

- delete an image if it does not clearly relate to the text immediately before or after it. For example, if the image shows the final dish while the text is about the ingredients and the first step of the recipe, remove the image.

2. Remove duplicate images:

- if multiple consecutive images show the same thing, keep only the first.

3. Reorder the images and text chunks:

- if an image clearly belongs with the following text chunk, you may swap the image with that next text chunk. No other rearranging is allowed.

Output:

Return a JSON object in the following format:

```
{
  "reasoning": {
    "image1": {"reasoning": "...", "operation": "keep/remove/replace"},
    ...
  },
  "introduction": [ ... processed introduction sequence ... ],
  "content": [ ... processed content sequence ... ]
}
```

Use "replace" only when swapping an image with the text chunk immediately after it.

Introduction:

```
[
  {input_introduction}
]
```

Content:

```
[
  {input_content}
]
```

After the steps described above, we obtain a cleaned interleaved sequence. We then prompt the VLM with the following instruction to convert this sequence into a user–assistant conversation.

Prompt for Generating Conversation from Interleaved Content (Qwen2.5 VL 72B)

You are a skilled data annotation expert tasked with converting raw documents into conversational training data for a chatbot. Each document contains interleaved text and images. Your goal is to transform this document into a single-turn, natural conversation between a human user and a chatbot. The conversation may take the form of question–answering, story-telling, or simply giving a list of facts or tips. For example, if the document is a recipe, the user might ask the chatbot how to make the dish, and the chatbot would respond step by step.

Input

The input consists of two parts, each as a list of interleaved text and images:

- **Introduction:** Includes the title and background context for the document. It may also include images (e.g., the final outcome of a tutorial). You should decide whether these images are relevant for the user's input.
- **Content:** The main body of the document. This contains the core instructions, steps, or detailed discussion points.

Task

Read both introduction and content, then convert them into a structured conversation following this format:

```
```json
[
 {
 "role": "user",
 "content": [
 {
 "type": "image",
 "image": "image1"
 },
 {
 "type": "text",
 "text": "How to make a custom wooden baking dish like this?"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "type": "text",
 "text": "Step 1: ..."
 },
 {
 "type": "image",
 "image": "image2"
 },
 ...
]
 }
]
```
```

Requirements

- Output must be valid JSON in the structure above.
- Do not change the order of images in the raw content.
- You need to paraphrase or rewrite text to make the dialogue natural and coherent (e.g., convert first-person narration into user–assistant exchanges).
- Always use a single round of conversation: the user appears once at the beginning, and the assistant appears once at the end.
- If images from the introduction are relevant to the user's query, include them in the user's input. If they are only loosely related (e.g., a butterfly image in a tutorial about making my backyard more attractive for butterflies), exclude them.
- Ensure the conversation feels realistic, contextually relevant, and faithful to the document's topic. ""

Introduction:

{input introduction}

Content:

{input content}

Note that in this step, only the images from the introduction section are fed to the VLM, allowing it to determine whether and where images should appear in the user's query. All images in the main content remain as placeholders when provided to the VLM.

7. Synthetic Interleaved Generation Data

To construct the prompts for general-purpose instruction-tuning data, we start from eight base categories: “*Sports*”, “*Outdoor & Survival*”, “*DIY & Crafting*”, “*Vehicle & Transportation*”, “*Personal Care & Health*”, “*Farm, Pet, and Animals*”, “*Home & Living*”, and “*Office & Productivity*”. We ask eight human annotators to further decompose these base categories into 151 subcategories and to write approximately 10 seed questions for each subcategory.

We then prompt OpenAI O3 [22] with the instruction shown below to expand these seed questions, yielding 15,270 questions in total. To reduce duplication and improve generation quality, we use the highest thinking budget and provide the model with the other subcategories and their seed questions within the same base category as additional context.

Prompt for Expanding the Seed Questions (O3)

I am building a how-to dataset for training a multimodal chatbot that can generate step-by-step guides or articles with interleaved text and images.

Given a category name and a subcategory name, along with some seed questions, your task is to generate additional how-to questions that fit the subcategory.

Follow these rules carefully:

Expand and specialize:

+ If any seed question can be broken down into more specific versions, do so.

+ For example, the question "How to perform CPR on an adult (and key differences for children)?" should be expanded into: "How to perform CPR on an adult?" and "How to perform CPR on a child?"

+ Every generated question must be clearly related to both the subcategory and its parent category.

+ Quantity and diversity: generate as many relevant and diverse questions as possible, covering different angles, contexts, and variations of the subcategory topic.

+ Avoid duplication or misclassification: I will also provide you with questions from other subcategories within the same category. Use these as reference to ensure you do not duplicate existing questions or generate ones that belong to another subcategory.

Input:

Category: *{category name}*

Subcategory: *{subcategory name}*

Seed questions:

{seed questions}

Questions from other subcategories:

{other subcategories under the same category and their seed questions}

Output format:

Return the generated questions as a JSON array:

```
[
  "question1",
  "question2",
  ...
]
```


8. Interleaved Context Alignment Data

Table A5 lists all datasets used in the interleaved context-alignment stage. We adopt open-source datasets spanning text-to-image generation, image editing, and multi-reference generation tasks.

| Data Source | Split | Task | Number (k) | Note |
|---------------------------|-------------------|--------------------------------------|--------------|-------------------------------|
| Open-GPT-4o-Image [5] | Generation | Text-to-Image | 39k | Generated by GPT-4o |
| BLIP3o [3] | Generation | Text-to-Image | 60k | Generated by GPT-4o |
| TextAtlas5M [27] | TextScenesHQ | Text-to-Image | 36k | Text rendering |
| TextAtlas5M [27] | Cover Book | Text-to-Image | 207k | Text rendering |
| TextAtlas5M [27] | LongWordsSubset-A | Text-to-Image | 266k | Text rendering |
| UniWorld-V1 [17] | osp | Text-to-Image | 286k | - |
| ShareGPT-4o-Image [2] | Generation | Text-to-Image | 45k | Generated by GPT-4o |
| Sum | | Text-to-Image | 939k | |
| OmniGen [33] | Subject-driven | Image Editing | 192k | Multi-reference generation |
| GPT-IMAGE-EDITE-1.5M [30] | - | Image Editing | 1500k | - |
| Open-GPT-4o-Image [5] | Editing | Image Editing | 41k | - |
| ShareGPT-4o-Image [2] | Editing | Image Editing | 46k | - |
| OmniGen2 [32] | Video Edit | Image Editing | 1083k | From video sequence |
| OmniGen2 [32] | Video ICEdit | Image Editing | 155k | From video sequence |
| OmniGen2 [32] | Video ICGen | Image Editing | 888k | From video sequence |
| SEED-Data-Edit [8] | - | Image Editing | 77k | - |
| StyleBooth [12] | - | Image Editing | 11k | - |
| NHR-Edit [13] | - | Image Editing | 358k | Generated by a in-house model |
| Echo-4o-Image [37] | - | Image Editing | 73k | - |
| Pico-Banana-400k [25] | - | Image Editing | 400k | - |
| NanoConsistent [37] | - | Image Editing | 150k | - |
| ImgEdit [38] | - | Image Editing | 184k | Filtered by VLM |
| UniWorld-V1 [17] | Ghibli | Image Editing | 36k | - |
| UniWorld-V1 [17] | ip_img | Image Editing | 24k | - |
| UniWorld-V1 [17] | omiedit | Image Editing | 368k | - |
| Sum | | Image Editing | 5586k | |
| OmniGen2 [32] | Video Interleaved | Interleaved Context Alignment | 657k | - |
| Video Dense Caption | - | Interleaved Context Alignment | 5000k | Introduced in Sec. ?? |
| Sum | | Interleaved Context Alignment | 5657k | |
| Interleaved Generation | robot planning | Interleaved Generation | 168k | Introduced in Sec. ?? |
| Interleaved Generation | cooking | Interleaved Generation | 15k | Introduced in Sec. ?? |
| Interleaved Generation | general VQA | Interleaved Generation | 15k | Introduced in Sec. ?? |
| Interleaved Generation | web pages | Interleaved Generation | 268k | Introduced in Sec. ?? |
| Sum | | Interleaved Generation | 466k | |

Table A5. Training data of the interleaved context alignment stage.

References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 23
- [2] Junying Chen, Zhenyang Cai, Pengcheng Chen, Shunian Chen, Ke Ji, Xidong Wang, Yunjin Yang, and Benyou Wang. Sharegpt-4o-image: Aligning multimodal models with gpt-4o-level image generation. *arXiv preprint arXiv:2506.18095*, 2025. 24
- [3] Jiu-hai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 14, 24
- [4] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 14
- [5] Zhihong Chen, Xuehai Bai, Yang Shi, Chaoyou Fu, Huanyu Zhang, Haotian Wang, Xiaoyan Sun, Zhang Zhang, Liang Wang, Yuanxing Zhang, et al. Opengpt-4o-image: A comprehensive dataset for advanced image generation and editing. *arXiv preprint arXiv:2509.24900*, 2025. 24
- [6] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 14
- [7] Yi Dong, Yusuke Muraoka, Scott Shi, and Yi Zhang. Mm-food-100k: A 100,000-sample multimodal food intelligence dataset with verifiable provenance. *arXiv preprint arXiv:2508.10429*, 2025. 23
- [8] Yuying Ge, Sijie Zhao, Chen Li, Yixiao Ge, and Ying Shan. Seed-data-edit technical report: A hybrid dataset for instructional image editing. *arXiv preprint arXiv:2405.04007*, 2024. 24
- [9] Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, and Ying Shan. Making llama see and draw with seed tokenizer. *arXiv preprint arXiv:2310.01218*, 2023. 13
- [10] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. In *NeurIPS*, 2023. 14
- [11] Google DeepMind. Gemini 2.5 flash image (“nano banana”) — image generation & editing model, Aug. 2025. Released August 26, 2025; state-of-the-art image generation and editing model. 13
- [12] Zhen Han, Chaojie Mao, Zeyinzi Jiang, Yulin Pan, and Jingfeng Zhang. Stylebooth: Image style editing with multimodal instruction. In *CVPR*, pages 1947–1957, 2025. 24
- [13] Maksim Kuprashevich, Grigorii Alekseenko, Irina Tolstykh, Georgii Fedorov, Bulat Suleimanov, Vladimir Dokholyan, and Aleksandr Gordeev. Nohumansrequired: Autonomous high-quality image editing triplet mining. *arXiv preprint arXiv:2507.14119*, 2025. 24
- [14] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 14
- [15] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, et al. Flux. 1 kontekst: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025. 14
- [16] Ang Li, Charles Wang, Kaiyu Yue, Zikui Cai, Ollie Liu, Deqing Fu, Peng Guo, Wang Bill Zhu, Vatsal Sharan, Robin Jia, et al. Zebra-cot: A dataset for interleaved vision language reasoning. *arXiv preprint arXiv:2507.16746*, 2025. 13
- [17] Bin Lin, Zongjian Li, Xinhua Cheng, Yuwei Niu, Yang Ye, Xianyi He, Shenghai Yuan, Wangbo Yu, Shaodong Wang, Yunyang Ge, et al. Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147*, 2025. 14, 24
- [18] Minqian Liu, Zhiyang Xu, Zihao Lin, Trevor Ashby, Joy Rimchala, Jiabin Zhang, and Lifu Huang. Holistic evaluation for interleaved text-and-image generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 22002–22016. Association for Computational Linguistics, 2024. 2
- [19] Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, Jiaqi Liao, Yingming Wang, Honghao Fu, Chunrui Han, et al. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*, 2025. 14

- [20] OpenAI. Dall-e 3. <https://openai.com/blog/dall-e-3>, 2023. 14
- [21] OpenAI. Gpt-4o: The next-generation omni model, May 2024. Introduces GPT-4o, a multimodal model capable of real-time audio, vision, and language. 14
- [22] OpenAI. Openai o3 and o4-mini system card. System card, OpenAI, Apr. 2025. 22
- [23] Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiahai Chen, Kunpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025. 14
- [24] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 14
- [25] Yusu Qian, Eli Bocek-Rivele, Liangchen Song, Jialing Tong, Yinfei Yang, Jiasen Lu, Wenze Hu, and Zhe Gan. Pico-banana-400k: A large-scale dataset for text-guided image editing. *arXiv preprint arXiv:2510.19808*, 2025. 24
- [26] Luozheng Qin, Jia Gong, Yuqing Sun, Tianjiao Li, Mengping Yang, Xiaomeng Yang, Chao Qu, Zhiyu Tan, and Hao Li. Uni-cot: Towards unified chain-of-thought reasoning across text and vision. *arXiv preprint arXiv:2508.05606*, 2025. 14
- [27] Alex Jinpeng Wang, Dongxing Mao, Jiawei Zhang, Weiming Han, Zhuobai Dong, Linjie Li, Yiqi Lin, Zhengyuan Yang, Libo Qin, Fuwei Zhang, et al. Textatlas5m: A large-scale dataset for dense text image generation. *arXiv preprint arXiv:2502.07870*, 2025. 24
- [28] Guo-Hua Wang, Shanshan Zhao, Xinjie Zhang, Liangfu Cao, Pengxin Zhan, Lunhao Duan, Shiyin Lu, Minghao Fu, Xiaohao Chen, Jianshan Zhao, et al. Ovis-u1 technical report. *arXiv preprint arXiv:2506.23044*, 2025. 14
- [29] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. 14
- [30] Yuhan Wang, Siwei Yang, Bingchen Zhao, Letian Zhang, Qing Liu, Yuyin Zhou, and Cihang Xie. GPT-Image-Edit-1.5M: A million-scale, gpt-generated image dataset. *arXiv preprint arXiv:2507.21033*, 2025. 24
- [31] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 14
- [32] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, et al. Omnigen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*, 2025. 14, 24
- [33] Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. In *CVPR*, pages 13294–13304, 2025. 14, 24
- [34] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. In *ICLR*, 2025. 14
- [35] Yi Xin, Qi Qin, Siqi Luo, Kaiwen Zhu, Juncheng Yan, Yan Tai, Jiayi Lei, Yuwen Cao, Keqi Wang, Yibin Wang, et al. Lumina-dimoo: An omni diffusion large language model for multi-modal generation and understanding. *arXiv preprint arXiv:2510.06308*, 2025. 14
- [36] Ling Yang, Ye Tian, Bowen Li, Xinchun Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada: Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025. 14
- [37] Junyan Ye, Dongzhi Jiang, Zihao Wang, Leqi Zhu, Zhenghao Hu, Zilong Huang, Jun He, Zhiyuan Yan, Jinghua Yu, Hongsheng Li, et al. Echo-4o: Harnessing the power of gpt-4o synthetic images for improved image generation. *arXiv preprint arXiv:2508.09987*, 2025. 24
- [38] Yang Ye, Xianyi He, Zongjian Li, Bin Lin, Shenghai Yuan, Zhiyuan Yan, Bohan Hou, and Li Yuan. Imgedit: A unified image editing dataset and benchmark. *arXiv preprint arXiv:2505.20275*, 2025. 14, 24
- [39] Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*, 2025. 14
- [40] Kaizhi Zheng, Xuehai He, and Xin Eric Wang. Minigt-5: Interleaved vision-and-language generation via generative vokens. *arXiv preprint arXiv:2310.02239*, 2023. 13