The future is vast.
What will you build?

Dreamers is a research, data science, and tech consultancy that specializes in building solutions for complex problems.

We are based in sunny California, but work remotely with clients from all over the world. We're engineer-owned and operated, meaning you'll get personal attention and quick answers without having to bushwhack through obfuscating layers of management.

Palazzo
[Case Study]

The Vision

Imagine walking into a room, snapping a picture, and instantly getting recommendations for better furniture. Not just generic suggestions, but real pieces from actual catalogs—perfectly sized, styled, and placed in your space. The goal of this project was to bring that experience to life, using a blend of cutting-edge visual analysis, smart recommendations, and photorealistic rendering.

The Challenge

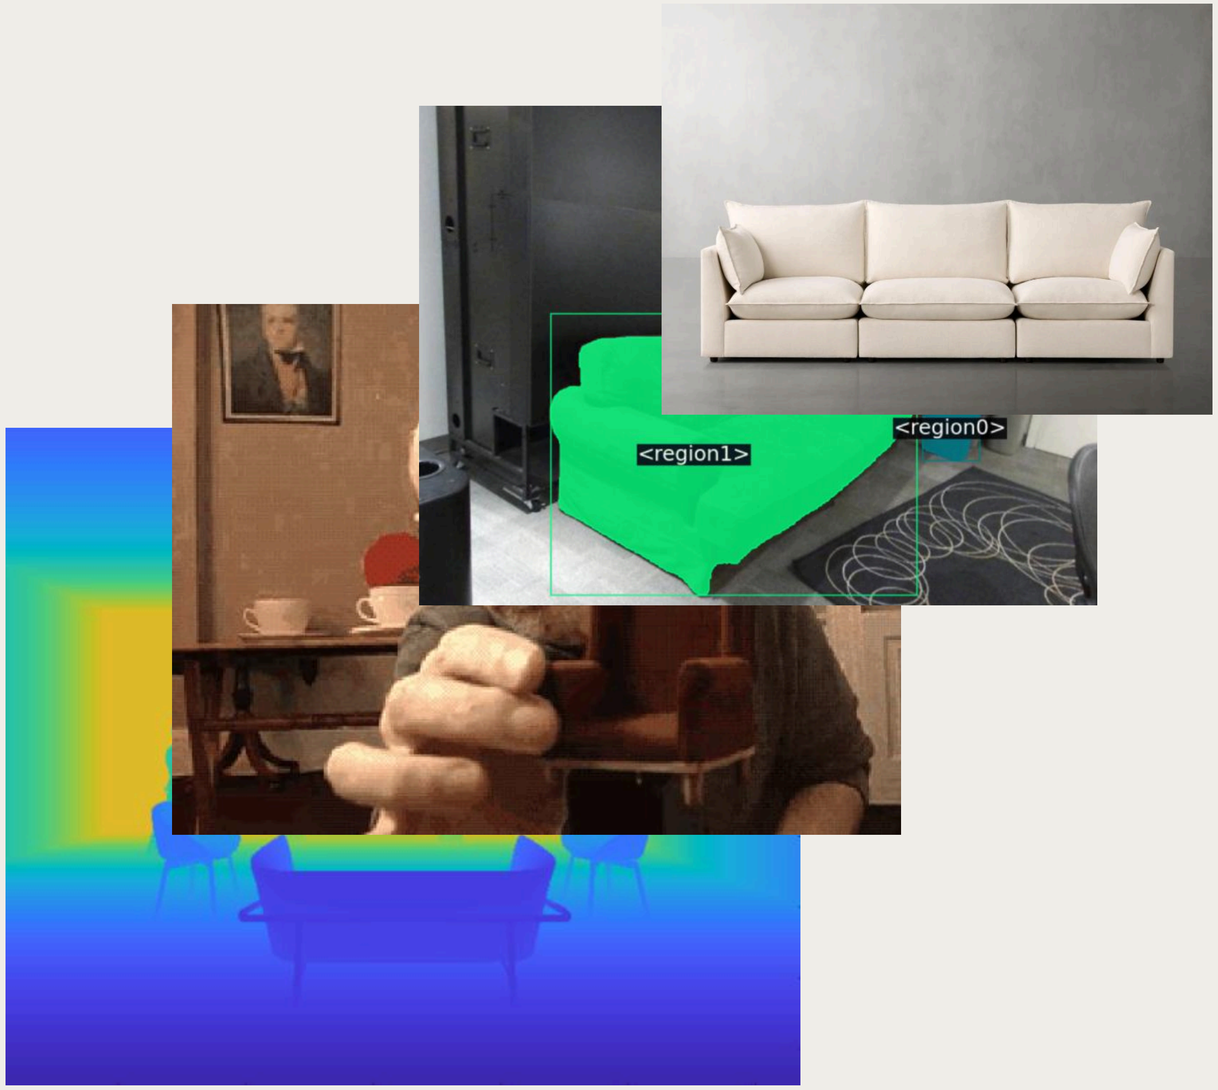

Creating a system capable of this level of design intelligence is no trivial task. It involves not just recognizing objects and spaces but truly understanding them—understanding depth, perspective, and scale while maintaining the style and aesthetic harmony of the room. Blending 2D catalog images into 3D environments, generating lifelike meshes, and ensuring realistic placement and lighting require a seamless integration of computer vision, machine learning, and 3D modeling. The complexity of building such a system lies in the delicate balance between technological precision and artistic presentation.

To make this vision a reality, we had to build a system that could:

- **Visualize the Space:** Develop a model that could analyze and configure the room layout.
- **Leverage Real Catalogs:** Create a robust database with data from leading furniture companies.
- **Make Smart Suggestions:** Implement a **recommender module** that not only matches style and color but also ensures a perfect fit in the room.
- **Transform 2D to 3D:** Generate 3D meshes from 2D images of catalog items.
- **Understanding 3D, Given a 2D Image:** Monocular depth perception—understanding depth from a single image—is technically impossible. Humans rely on binocular vision, using the slight differences between two eye perspectives to gauge depth. To overcome this, we used advanced machine learning solutions to generate a reliable depth map from 2D images, enabling the system to understand spatial relationships in the room.
- **Perfect the Fit:** Assess the pose and scale of objects being replaced, then rotate and fit new items seamlessly.
- **Deliver Realism:** Blend and stitch the new objects into the room with a photorealistic touch.

Obstacles We Faced

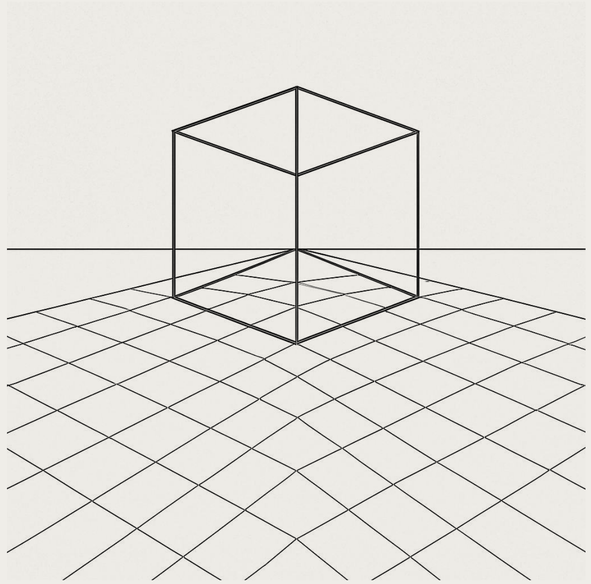

- **Ridiculously Tight Timeline:** Solution: Work like maniacs. Stay laser-focused and crank out results.
- **No Hosting from Palazzo:** Solution: We hosted the entire project on our Dreamers server, from development to production, avoiding delays and delivering on time.
- **Slow and Expensive 3D Mesh Tools:** Solution: We built our own. Using state-of-the-art tools, we followed top repositories and connected with the brightest minds. We remained aware of the daily shifts in cutting-edge technology, maintaining close communication with the three top teams at the forefront of 3D visual systems—including groups working on NERF, and specialized researchers at UCSD, KAUST, and NVIDIA. Our in-house solution was not only free and fast but also outperformed commercial options when the project began.
- **Evolving Technology Landscape:** The pace of change in this field is staggering. The tools available when we started were nearly obsolete by the time we finished. We adapted continuously, evolving our approach as new research and tools emerged. Modern problems require this level of dynamic reaction to a changing field. Anything less risks delivering outdated solutions by the time they're implemented.
- **Missing Room Dimensions:** Originally, room dimensions were supposed to be provided as part of the project scope. When it became clear they would not be, we built a system to estimate room dimensions accurately, keeping us on track despite a dynamic and evolving scope.
- **Performance Bottlenecks:** Solution: The model we needed for scale and orientation estimation was powerful but slow (300 seconds per run). That wasn’t going to fly. We brought in our high-performance computing expert and slashed processing time to 10 seconds.

The process
[Under the hood]

Why Depth Perception is Hard, Especially with A Single Image

When an image is captured using a single lens (monocular vision), it lacks the explicit depth information humans perceive using two eyes and stereoscopic vision. So, to create new 3D from an image, we must infer spatial relationships from shading, perspective, and object size. Even then, there is inherent ambiguity in some situations, as demonstrated by the gif below. This makes tasks like accurately replacing furniture and reconstructing 3D scenes incredibly challenging.

See for Yourself: https://cdn-uploads.huggingface.co/production/uploads/660710b03ef451aa2bab8971/_Pflg6HQGjKpyVGDC9MBd.gif

In order to address this problem, we build a **Depth Map** of the space. A depth map is a representation where each pixel encodes depth information, typically using grayscale values—the darker the pixel, the closer the object; the lighter, the farther away. Depth maps allow AI systems to estimate spatial relationships, helping to reconstruct a 3D scene from flat images.

Below is a depth map demonstrating how different objects in a scene are positioned relative to the camera.

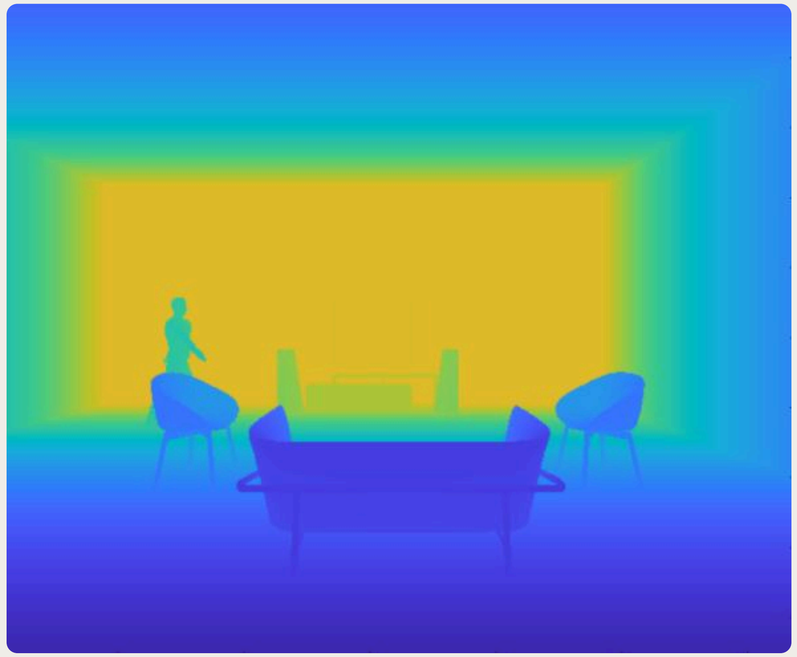

To successfully attack the furniture modification problem, we needed an **Image Classifier** capable of knowing exactly what every item in the room is, and which pixels map to that item. This is known as classification and masking. Here is an example of **Masking**, in which the pixels map appropriately to what we want to replace in the room.

In addition to “masking” images, we want to know what each item is, and draw a **3D Bounded Box** around it, so that the computer vision system can play around in 3 dimensions.

The System in Action:

Example 1:

Prompt: **“Our kids playroom has a sofa that can really use an upgrade. I like the color, but I want something more modern and clean looking, roughly the same size.”**

Supplied Photo:

Sofa selected by database

System Solution:

Selected sofa has been Sized, Aligned, and Blended with appropriate lighting to fit into the room.

Example 2:

Prompt: “The L shape in my living area takes up a lot of room. Can we replace it with a sofa and a few other pieces to let the space breathe? Pick colors that you feel go well with the paint.”

Supplied Photo:

Sofa selected by database

Looks Like a Couch in a Room - Isn't This Problem Easy?

At first glance, replacing a couch in a room seems trivial. With all our modern technology, why isn't this just a drag-and-drop process? The difficulty lies in the fact that the input image we receive has different angles—both vertically and horizontally—than the reference catalog image. Even defining what the original angle is, on its own, is not an easy problem. This means we can't just place the new furniture in the scene; we need to generate a completely new image that correctly matches the perspective of the room. Further, lighting in the room rarely matches the lighting from the original product image. Shadows, reflections, and color tones all have to be adjusted dynamically, or else the replacement will feel unnatural. Additionally, if the new sofa is placed at an angle but its size doesn’t subtly decrease with distance, the human brain instantly perceives something is very off, even without being able to pinpoint what that is.

The fact that these modifications appear seamless is a significant achievement, not something that works automatically out of the box.

Results

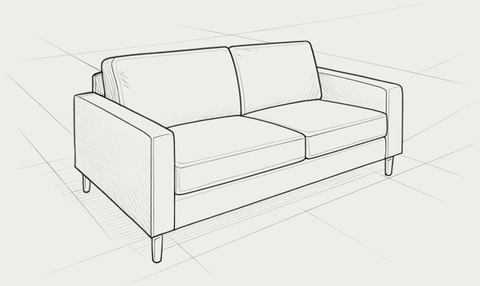

Despite the hurdles, we delivered a fast, high-quality solution that has set a new benchmark for interior design automation. The project exceeded Palazzo's expectations, combining advanced technology with creative problem-solving to turn an ambitious vision into a functional reality. In turn, their clients, including household names such as Decorator's Best, Ashley, and Arhaus, will now be able to better present their furniture lines to the buying public.

"I loved the founder's vision and technical understanding."

For this and other rave reviews, you may read about us here:

https://clutch.co/go-to-review/0533cf36-1cd4-4c33-9eab-5bc7d9fdf501/361001

Conclusion

The way humans interface with computers is evolving rapidly. Search is dead. We're moving beyond rigidly confined queries and into an era where we interact with data systems using our own words, where natural language becomes the primary interface. Retail is transforming, driven by deeply personalized compute resources and intelligent systems that understand context and intent. We're also starting to see hints of augmented reality creeping into our visual spaces, blending the physical and digital worlds seamlessly. **Dreamers Inc. is uniquely positioned to accept this challenge and push the boundaries.**

The future is vast.
What will you build?