

A Jalgorithm for Japplying Jeans to Jobjects

Dave Pagurek

March 2023

Abstract

For years, society has been plagued by the lack of a conclusive answer to the question of how arbitrary things would wear jeans, if they wore jeans. This paper proposes a jalgorithm to determine the mathematically optimal configuration of pants on a silhouette, leveraging recent insights in the fields of computer vision and differentiable rendering.

1 Introduction

In recent years, the internet has been ruminating over the question of how various objects would wear pants, if they hypothetically wore pants. [2] It began in 2015 using a dog as the object in question, [3] with other objects such as humans on their hands and knees [12] or pants themselves [4] being discussed later.

Despite being in the public discourse for nearly a decade, no conclusions have been reached. The authors believe this is due to a lack of a principled approach to the study of pants-wearing. This paper aims to rectify the situation by proposing a formula to score the quality of the application of pants to an object, and writing software to optimally apply pants under this system.

2 Background

Gradient-based optimization has been a subject of study in mathematics for a long time. The core idea is that if one has an objective function one wishes to maximize or minimize, one can use the gradient of the output with respect to the function's inputs to determine how to adjust the inputs in order to move the output closer to the desired outcome. There are many algorithms, including Adam optimization [5], to update inputs in this manner using gradient information.

Differentiable rendering is a technique in the world of computer graphics and computer vision that enables gradient-based optimization of functions involving rendering of images. If the rendering process is itself differentiable, then one could, for example, optimize the representation of a 3D scene such that it matches a reference photo when rendered. [7, 8] The diffvg library [6] provides a differentiable vector graphics renderer, allowing one to optimize 2D vector shapes such as lines, curves, and polygons to produce desired visuals.

3 Method

3.1 Algorithm

Our pants-fitting jalgorithm takes as input a vector shape representing a silhouette on which one wants



Figure 1: Prototypical jeans: \mathbf{X} , the initial pants shape we apply to objects.

to apply pants. We rasterize this to a 256×256 -pixel image, concerning ourselves only with the alpha channel. We will refer to this one-channel silhouette image as the 256×256 matrix S .

We then create a polygon representing the pants we want to fit onto the silhouette, whose vertex positions X we will optimize. We initialize its vertices to \mathbf{X} , the shape of prototypical jeans, shown in Figure 1. We then solve for the vertices \hat{X} that minimize pants-fitting loss, $l(X)$, described in detail in Section 3.2, using Adam optimization [5] and gradients provided via diffvg [6] vector rendering and Torch [11] for 100 iterations:

$$\hat{X} = \operatorname{argmin}_X l(X) \quad (1)$$

3.2 Loss Function

Our loss function is motivated by insights from the authors’ multiple decades of pants-wearing experience. We propose the following rules, which we then will encode mathematically:

1. **Duality of man:** Pants should bisect the silhouette of the wearer through its center, covering half its area.
2. **No-stretch jeans:** Pants, when applied to the wearer in rest position, should not be so deformed that they no longer resemble prototypical jeans.
3. **(Topologically) torn jean avoidance:** Limbs should only go through the open ends of pants, and similarly, closed ends of pants should not intersect limbs.
4. **Dress code compliance:** Pants should be on the body of the wearer, not off of it.

From Rule 1, we get the following loss terms, where x_{toleft} and x_{topright} correspond to the endpoints of the waistband of prototypical jeans, and $J(X)$ is the 256×256 matrix of the rendered alpha channel of the jeans vertices X . L_{center} incentivizes waistbands that go through the center of the silhouette, and L_{coverage} incentivizes jeans that cover half of the silhouette.

$$L_{\text{center}} = \left\| \frac{x_{\text{toleft}} + x_{\text{topright}}}{2} - \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \right\|^2 \quad (2)$$

$$L_{\text{coverage}} = \left(\frac{1}{2} - \frac{\sum_{i,j} \max((S - J(X))_{ij}, 0)}{\sum_{i,j} S_{ij}} \right)^2 \quad (3)$$

From Rule 2, we get the loss term L_{rigid} , which incentivizes the jean points X to be as close as possible to an affine transformation of the prototypical jean points \mathbf{X} by finding the error of the least-squares-minimal transformation T that describes the relation between X and \mathbf{X} .

$$L_{\text{rigid}} = \|X - T\mathbf{X}\|^2, \quad (4)$$

$$T = \begin{bmatrix} \hat{p}_1 & \hat{p}_2 & \hat{p}_3 \\ \hat{p}_4 & \hat{p}_5 & \hat{p}_6 \\ 0 & 0 & 1 \end{bmatrix}, \quad (5)$$

$$\hat{p} = \operatorname{argmin}_p \left\| \begin{bmatrix} \mathbf{X}_{0,1} \\ \mathbf{X}_{1,1} \\ \vdots \\ \mathbf{X}_{0,n} \\ \mathbf{X}_{1,n} \end{bmatrix} - \begin{bmatrix} X_{0:1,1} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & X_{0:1,1} & 1 \\ \vdots & & & & \\ X_{0:1,n} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & X_{0:1,n} & 1 \end{bmatrix} p \right\|^2 \quad (6)$$

From Rule 3, we get the terms L_{open} and L_{closed} , terms incentivizing open and closed edges of the jeans respectively to and to not intersect the silhouette, and L_{inter} , incentivizing edges to not intersect each other. We refer to the set of all edges as $E(X)$, and the of open and closed edges as $E_o(X)$ and $E_c(X)$. The rendered alpha channel of an edge with stroke thickness of 5 is referred to as $R(e)$.

$$L_{\text{open}} = \frac{-\sum_{i,j} R(E_o(X))_{ij} S_{ij}}{\sum_{i,j} R(E_o(X))_{ij}} \quad (7)$$

$$L_{\text{closed}} = \frac{\sum_{i,j} R(E_c(X))_{ij} S_{ij}}{\sum_{i,j} R(E_c(X))_{ij}} \quad (8)$$

$$L_{\text{inter}} = \sum_{\{e_1, e_2\}: \binom{E(X)}{2}} \frac{\sum_{i,j} R(e_1)_{ij} R(e_2)_{ij}}{\sum_{i,j} \min(R(e_1)_{ij} + R(e_2)_{ij}, 1)} \quad (9)$$

From Rule 4, we get L_{inside} , discouraging the pants from being off the silhouette:

$$L_{\text{inside}} = \frac{\sum_{i,j} \max(J(X)_{ij} - S_{ij}, 0)}{\sum_{i,j} J(X)_{ij}} \quad (10)$$

We take a weighted sum over all the loss terms to get the final loss function:

$$l(X) = \sum_i k_i L_i \quad (11)$$

For our results, we used the following weights:

$$\begin{aligned} k_{\text{center}} &= 50, \\ k_{\text{coverage}} &= 100, \\ k_{\text{rigid}} &= 5, \\ k_{\text{open}} &= 1, \\ k_{\text{closed}} &= 1, \\ k_{\text{inter}} &= 50, \\ k_{\text{inside}} &= 1 \end{aligned}$$

4 Results and Analysis

4.1 Validation

To verify that the algorithm outlined in Section 3 operates as expected for known cases, we ran it on the



Figure 2: Placement of jeans on a human via the algorithm, matching conventional wisdom of where jeans belong on a human.




Figure 3: Placement of jeans on the Vitruvian Man, the body shape we all aspire to have.

silhouette of a human. Figure 2 shows that the algorithm placed the jeans on the legs of the human, which the reader may recall is the standard configuration of human jeans.

While this is only a sample size of one, there is in fact only one species for whom jeans are originally designed, so we feel confident in our ability to extrapolate from this to other figures. However, as a gift to the reader, we test on one more human-shaped validation input: Leonardo da Vinci's *Vitruvian Man*. Conceived as a depiction of ideal body proportions, [1] the octopedal body plan represents what we all aspire to look like. Figure 3 shows the algorithm confidently placing jeans on two of the four legs, fashionably rolling up one leg in a manner reminiscent to one-strapping a backpack. [9]




Figure 4: Placement of jeans on a dog via the algorithm. The jeans are placed on all legs, but in an interesting deviation from the original hypothesis, the jeans leg for the rear legs has ridden up to align with the derriere.

4.2 Novel Inputs

We next tested the algorithm on a dog in an attempt to settle the original debate. The original image presented two hypotheses for dog jean placement: on the rear legs only, subdividing the dog into two horizontal sections, or on all legs, subdividing the dog into two vertical sections. [3] Figure 4 shows that the algorithm places the jeans in a configuration much more like the latter than the former, but with a twist: rather than form-fitting the rear legs, the algorithm chose to collapse the rear pants leg so that it ends at the derriere of the dog.

In Figure 5, we present jeans applied to some more body shapes. First, one best described as horse-inspired (the literature shows that one likely cannot draw a good horse without reference material present [10]) dons a single pant leg. Here, the algorithm demonstrates that it would rather shrivel up a leg into nothing than stretch a waistband across any part of a horse. Next, a three-legged alien wears jeans on one and a half legs, with a small branch of jean reaching up to its shoulder for support. This demonstrates the algorithm's ability to balance its creative approach to jean design with functional concerns, like extra shoulder support. Finally, a spiky ball dragon wears jeans on only one of its ball legs, further demonstrating a preference for one-strapping.




Figure 5: Jean placements on alternative body shapes. Note the creative one-strap overall look of the tripod, and the bold asymmetric style of covering one ball of the ball dragon.

5 Conclusion

We have presented a mathematically sound optimization approach to determining the optimal configuration of pants on any object, should it choose to wear pants. Now that a definitive answer exists to this question, we expect no further arguments to be made on the subject on the internet. Since the algorithm’s outputs are, by definition, optimal, one may choose to take lessons from its output, and begin casually rolling up just one pant leg in one’s outfits.

References

- [1] Vitruvian man. *Wikipedia*. URL: https://en.wikipedia.org/wiki/Vitruvian_Man.
- [2] If a dog wore pants. *Know Your Meme*, Dec 2015. URL: <https://knowyourmeme.com/memes/if-a-dog-wore-pants>.
- [3] If a dog wore pants would he wear them like this or like this? *Utopian Raspberry - Modern Oasis Machine*, Dec 2015. URL: <https://www.facebook.com/1ur.mom/photos/a.535344933281912.1073741828.535340893282316/604404026376002/?type=3&theater>.
- [4] @bobby. quick question. *Twitter*, April 2017. URL: <https://twitter.com/bobby/status/849491087005810688>.
- [5] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL: <http://arxiv.org/abs/1412.6980>.
- [6] Tzu-Mao Li, Michal Lukáč, Gharbi Michaël, and Jonathan Ragan-Kelley. Differentiable vector graphics rasterization for editing and learning. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 39(6):193:1–193:15, 2020.
- [7] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019.
- [8] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. A general differentiable mesh renderer for image-based 3d reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [9] Phil Lord and Christopher Miller. 21 Jump Street. 2012.
- [10] Dave Pagurek. Can you draw a horse without a reference? Jul 2022. URL: <https://davepagurek.github.io/horse-drawings/>.
- [11] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [12] /u/flooid. If a man wore pants... *Reddit*, Dec 2015. URL: https://www.reddit.com/r/funny/comments/3yosk1/if_a_man_wore_pants/.