Operational Evidence and Outcomes
Volume 1 · Number 1 · 2026

More Recoverable Variation, Better Identity Retrieval? A Face Subspace Fitted on Other People

Christopher Harrison[1] · Liam Morgan[2] · Benjamin Lewis[3]

[1] Greece · [2] Estonia · [3] Latvia

Original publication: November 8, 2026 | Article 581

Abstract

A larger image subspace necessarily reconstructs more variation, but its usefulness for identity retrieval remains an empirical question. We fitted nested principal-component representations on 200 Olivetti photographs from 20 people, then used a separate 20-person gallery and 100 queries. Increasing the representation from eight to 32 coordinates raised nearest-gallery correctness from 79% to 93%; fifteen queries improved and one deteriorated. Mean reconstruction residual energy also fell, by 0.11960 on centered unit-norm inputs. These findings concern one externally fitted linear representation and one historical collection. They do not establish identity preservation in generated video.

Executed analysis of historical public photographs. Representation fitting and retrieval identities are disjoint; the retrieval task uses a known roster.

More Recoverable Variation, Better Identity Retrieval? A Face Subspace Fitted on Other People

> Executed analysis of historical public photographs. Representation fitting and retrieval identities are disjoint; the retrieval task uses a known roster.

Abstract

A larger image subspace necessarily reconstructs more variation, but its usefulness for identity retrieval remains an empirical question. We fitted nested principal-component representations on 200 Olivetti photographs from 20 people, then used a separate 20-person gallery and 100 queries. Increasing the representation from eight to 32 coordinates raised nearest-gallery correctness from 79% to 93%; fifteen queries improved and one deteriorated. Mean reconstruction residual energy also fell, by 0.11960 on centered unit-norm inputs. These findings concern one externally fitted linear representation and one historical collection. They do not establish identity preservation in generated video.

1. Two meanings of keeping more facial information

A representation can preserve visible variation without preserving the distinctions required to recognize a person. A direction with substantial variance may describe illumination or expression. A smaller direction may separate two identities. Increasing representation dimension changes both kinds of information together. Consequently, a claim about image reconstruction should be accompanied by a separate assessment of the identity decision that the representation will support.

Meng et al., 2026 introduces stacked multi-view identity mosaic injection for subject-preserving video generation. Its concern with retaining a subject across generated appearances motivates examining the evidence behind a representation's identity claim. The present experiment addresses a narrower preliminary question: when a face basis is fitted on other people, does increasing its dimension improve both reconstruction of held queries and retrieval of their identities? No video generator, mosaic injection, or learned identity encoder is executed.

We use a linear setting because the relationship between information retention and reconstruction is explicit. Principal components provide nested orthogonal subspaces: the first eight directions are contained in the first 32. The additional directions cannot increase orthogonal reconstruction loss for a fixed image. They can change nearest-neighbor assignments in either direction, however. The study therefore separates a mathematical consequence of projection from a measured change in recognition.

The split also gives representation transfer a precise meaning. Twenty identities provide every photograph used to fit the mean and basis. Another twenty supply the gallery and queries. Recognition concerns this second, known roster; it does not attempt to assign the names of the source identities. This arrangement prevents target photographs from choosing the basis while allowing genuine same-identity matches within the retrieval task.

2. Representation objectives and the identity question

Pearson, 1901 supplies the classical closest-subspace perspective behind the reconstruction objective. Tipping and Bishop, 1999 develops a probabilistic interpretation of principal-component analysis. Neither perspective makes retained variance an automatic measure of identity usefulness. We use deterministic projection rather than fitting a probabilistic observation model, so residual energy is a geometric quantity rather than a calibrated uncertainty estimate.

Eigenface recognition demonstrates the relevance of linear face representations to identification (Turk and Pentland, 1991). The contrast between eigenfaces and class-specific Fisher projections further clarifies that reconstruction and class discrimination are different fitting goals (Belhumeur et al., 1997). Our source fit is deliberately unsupervised with respect to identity: labels define the split, but are absent from the singular-value decomposition. This removes an opportunity to optimize directions directly for the twenty target classes.

Modern embedding methods provide useful boundaries for interpretation. FaceNet learns an embedding for recognition and clustering (Schroff et al., 2015), while ArcFace and CosFace organize discriminative learning through margin-based objectives (Deng et al., 2018; Wang et al., 2018). Our coordinates have neither their training objectives nor their learned invariances. A comparison of two PCA ranks therefore cannot rank these recognition systems or establish that a simple pixel baseline is sufficient for current applications.

Quality-aware face representations make another distinction relevant: a feature norm or representation dimension need not carry the same information as image suitability. AdaFace and MagFace explicitly address quality in their representation designs (Kim et al., 2022; Meng et al., 2021). We do not borrow their quality scores. Instead, reconstruction energy and correct identification remain separately measured outcomes, with no assumption that one can stand in for the other.

The first analytical proposition is directional and deterministic: using the larger nested subspace will weakly reduce query reconstruction energy. The second is empirical: the larger representation may improve the correctness of nearest-gallery identification on people excluded from basis fitting. The latter could fail if extra appearance variation changes distances unfavorably. We preserve individual reversals even when the average favors more coordinates.

3. Photographs, roles, and measurements

The analysis reads the captured Olivetti/AT&T MATLAB photograph resource (AT&T Laboratories Cambridge, 2026), with the historical face-recognition collection documented by Samaria and Harter, 1994. It contains 400 grayscale images with 4,096 values per image, organized as ten photographs for each of forty ordinal identities. These identity numbers support within-collection matching; they are not demographic labels, independent population samples, or evidence about a person's identity outside the collection.

Each image is centered by its own pixel mean and divided by its own Euclidean norm after centering. This makes the input insensitive to a common additive brightness change and positive multiplicative scale for that image. It does not remove all lighting effects, nonlinear intensity changes, pose differences, or occlusion. The normalized input has unit squared norm; the reconstruction outcome is the sum of squared residual coordinates, not mean squared error in original grayscale units.

A fixed seed, 20261108, permutes the forty identities. All 200 images from the first twenty identities fit the source representation. A separate seeded permutation within each of the remaining identities assigns five images to the gallery and five to queries. The target gallery thus has 100 photographs, and the held-query set has 100 photographs. Every query identity appears in the gallery, with five possible same-person references.

The source-image mean is computed once and subtracted from source, gallery, and query vectors. A thin singular-value decomposition of the 200 centered source vectors produces the common ordered basis. Rank eight and rank 32 use prefixes of exactly this basis, without whitening or a second normalization of

projected coordinates. No target image contributes to estimating the center or directions, and no query outcome chooses a rank.

Numerical arrays and linear algebra use the local NumPy and SciPy installations (Harris et al., 2020; Virtanen et al., 2020). Orthogonality and roundoff matter when comparing nested projections, as emphasized in numerical linear algebra (Higham, 2002) and floating-point analysis (Goldberg, 1991). Computation retains floating precision rather than rounding features to a storage format. The experiment therefore measures a coordinate budget, not an implemented compressed bitstream or inference-speed gain.

For each rank, the query is assigned the identity of its closest gallery photograph under squared Euclidean distance in projected coordinates. The first outcome is an indicator of a correct identity assignment. The second is reconstruction residual energy after mapping query coordinates back through the same basis and adding the source mean. Both outcomes use identical query rows, making within-query differences directly interpretable.

4. Model and conditional uncertainty

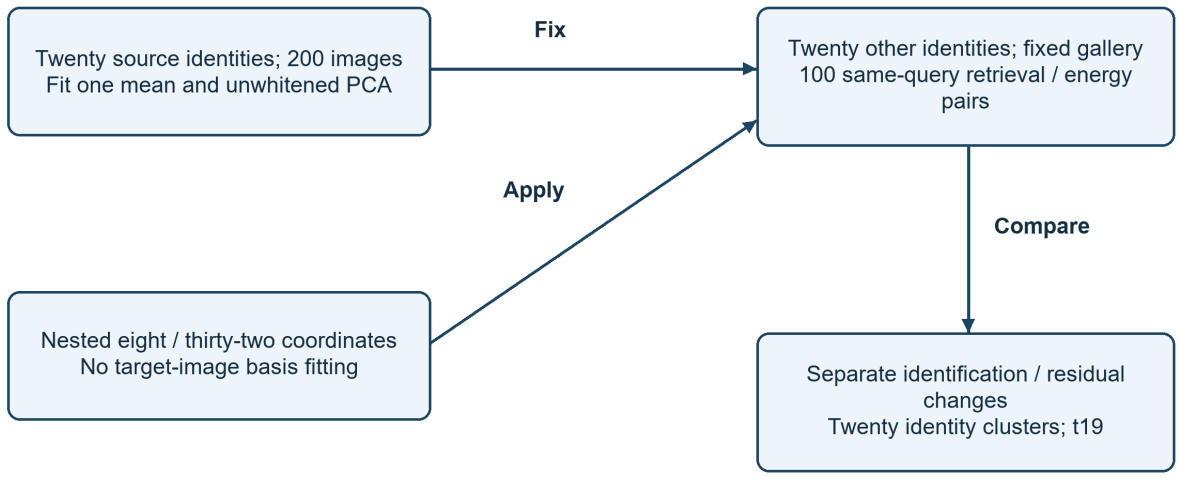

Figure 1. *Other identities fit the shared basis; nested coordinate budgets feed separate reconstruction and retrieval outcomes. Arrows describe analysis operations, not causal effects of an independently randomized representation treatment.*

We estimate two intercept-only regressions on paired query differences. Model 1 uses rank-32 correctness minus rank-eight correctness. Its intercept is the change in the proportion correctly identified. Model 2 uses the corresponding difference in reconstruction residual energy. Each query has equal weight; because every identity contributes five queries, this also gives equal total weight to each target identity.

The covariance estimator clusters the 100 rows by their twenty target identities, uses a finite-cluster correction, and refers confidence intervals to a t distribution with nineteen degrees of freedom. Repeated queries of the same person should not be treated as unrelated observations. The cluster rationale follows the dependence concerns discussed by Cameron and Miller, 2015, while White, 1980 motivates attention to nonconstant variation in regression errors. Neither reference makes twenty clusters a large or representative identity sample.

Both models condition on the realized source fit and common gallery. They omit variation from resampling source identities, refitting the basis, or changing which gallery photographs are retained. All identities also share the same fitted representation. The intervals are therefore working conditional summaries, rather than a complete uncertainty analysis for deploying a freshly trained representation in another collection. Intercept-only R-squared is reported for completeness and provides no useful discrimination diagnostic here.

Table 1. *Paired mean changes when increasing rank from eight to 32. Both models contain an intercept only, with no covariate adjustment; N=100 queries and twenty identity clusters. Standard errors are identity-clustered, and intervals use t19. Correctness is a proportion, while residual energy is dimensionless squared L2 on normalized inputs.*

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	0.1400 (0.0413)	[0.0536, 0.2264]	100
M1: fit	R-squared = 0.0000	-	-	100
M2	Mean outcome/contrast	-0.1196 (0.0085)	[-0.1373, -0.1019]	100
M2: fit	R-squared = 0.0000	-	-	100

Model labels: M1 = Rank-32 minus rank-8 correct-identification proportion; M2 = Rank-32 minus rank-8 query reconstruction residual energy.

5. The measured gains and the exception

Rank eight correctly identifies 79 queries, and rank 32 correctly identifies 93. Model 1 estimates a 0.14000 increase, or fourteen percentage points, with standard error 0.04129 and a 95% interval from 0.05357 to 0.22643. The interval describes this paired comparison under the stated covariance assumptions; it does not quantify uncertainty about the performance of an arbitrary future face encoder.

The transition table is more informative than the net gain alone. Seventy-eight queries are correct at both ranks, six are incorrect at both, fifteen change from incorrect to correct, and one changes from correct to incorrect. Thus the larger subspace does not dominate every decision. The adverse transition is retained because it shows exactly where a reconstruction-based expectation cannot guarantee individual recognition behavior.

Mean query reconstruction residual energy declines from 0.357258 at rank eight to 0.237659 at rank 32. Model 2 estimates a change of -0.119599, with standard error 0.008452 and interval [-0.137290, -0.101909]. On the source fitting set, the corresponding fractions of variation retained are approximately 55.28% and 81.68%. Source retained variance and target residual energy have different denominators and should not be presented as interchangeable percentages.

The negative reconstruction contrast is expected from nesting. Its magnitude describes how much target variation the added directions recover in this split. The positive correctness contrast is the separate empirical result. It would be incorrect to treat the small reconstruction p-value as proof that recognition must improve: the two regressions answer different questions, and only the first directly measures identity decisions.

There is also a useful distinction between the center used for reconstruction and the center's role in retrieval distances. Both a gallery image and a query subtract the same fitted source mean before projection. That common translation cancels when their projected coordinates are differenced. The fitted directions still matter to the distance, while the center remains necessary when interpreting the

reconstructed image. Documenting both operations avoids attributing an identity gain to a target-specific mean that was never estimated.

The larger representation uses four times as many coordinates per image, assuming the same numeric type for every coordinate. This arithmetic describes a feature-array budget only. A practical retrieval system would also store the fitted basis, source mean, gallery labels, and indexing structures, and might spend additional work evaluating distances. No such total-memory or latency comparison was measured here. The observed recognition gain should therefore be weighed against a declared coordinate count, rather than converted into an unsupported efficiency claim.

Method-comparison reasoning encourages inspecting disagreement rather than only average correspondence (Bland and Altman, 1986). Likewise, concordance concepts concern agreement on a declared scale (Lin, 1989). Here the outcomes are deliberately on different scales, so we do not calculate a single agreement coefficient between residual energy and a binary identity label. Keeping the units distinct follows the measurement concerns of Stevens, 1946 and prevents a lower reconstruction number from becoming an undefined general quality claim.

6. What this transfer experiment leaves unresolved

The historical collection is constrained relative to larger face benchmarks. LFW addresses faces in a less controlled setting (Huang et al., 2008), MegaFace examines recognition at scale (Kemelmacher-Shlizerman et al., 2015), and VGGFace2 emphasizes variation across pose and age (Cao et al., 2017). Those resources identify dimensions of external validity that are absent from this experiment. We do not use their images, add them to the effective sample size, or infer comparable performance from our observed 93%.

FairFace explicitly foregrounds demographic balance in dataset construction (Kärkkäinen and Joo, 2019). This experiment lacks the labels and sampling design required for an analogous assessment. A successful ordinal identity match does not establish demographic fairness, consent for new uses, or unrestricted permission to redistribute the source photographs. These issues remain separate from the numerical reproducibility of the local calculation.

The identity-disjoint basis fit addresses one possible source of overly favorable evaluation, but it does not resolve every dependence or selection problem. Kapoor and Narayanan, 2023 motivates tracing which information enters fitting. Roberts et al., 2017 emphasizes that validation structure should match the intended transfer question, and Varma and Simon, 2006 warns against using evaluation outcomes to choose a model. Our frozen two-rank comparison avoids test-driven rank selection, while remaining one exploratory comparison rather than a tuned model-selection exercise.

For subject-preserving generation, additional evidence would require generated appearances, reference images, detection availability, and an explicit rule for aggregating frames. A still-image gallery contains no temporal ordering, so its accuracy cannot demonstrate continuity through occlusion or motion. The present result is useful as an audit pattern: fit the representation outside the target identities, name the retrieval task, and report reconstruction and identity outcomes separately.

Dataset documentation should preserve the source roles, transformations, and unresolved permissions (Gebru et al., 2021). Reusable scientific records should also make the split and derived values findable and interpretable (Wilkinson et al., 2016). Accordingly, saved artifacts include the exact source and target identity lists, gallery and query row numbers, paired outcomes, model estimates, and basis-retained variation. Reconstructing the computation is possible without describing the fitted coordinates as a modern identity-preservation system.

7. Answer to the research question

Under an externally fitted face basis, the larger coordinate budget improves mean query reconstruction and known-roster identity retrieval in this captured split. The recognition gain comes from fifteen corrected assignments against one newly incorrect assignment. That exception matters: retained image variation has a guaranteed geometric relationship with projection loss, but its relationship with identity decisions remains empirical. The result supports reporting both quantities when evaluating a representation budget, with the fitting population and retrieval population stated explicitly.

References

1. Deng, Jiankang, et al. "ArcFace: Additive Angular Margin Loss for Deep Face Recognition" *arXiv preprint arXiv:1801.07698*, 2018. https://arxiv.org/abs/1801.07698.
2. White, Halbert. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. https://doi.org/10.2307/1912934.
3. Schroff, Florian, et al. "FaceNet: A Unified Embedding for Face Recognition and Clustering" *arXiv preprint arXiv:1503.03832*, 2015. https://arxiv.org/abs/1503.03832.
4. Roberts, David R., et al. “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure.” *Ecography*, vol. 40, no. 8, 2017, pp. 913-929. https://doi.org/10.1111/ecog.02881.
5. Samaria, F.S.; Harter, A.C.. (1994). Parameterisation of a stochastic model for human face identification. Proceedings of 1994 IEEE Workshop on Applications of Computer Vision. https://doi.org/10.1109/acv.1994.341300.
6. Wang, Hao, et al. "CosFace: Large Margin Cosine Loss for Deep Face Recognition" *arXiv preprint arXiv:1801.09414*, 2018. https://arxiv.org/abs/1801.09414.
7. Kärkkäinen, Kimmo and Jungseock Joo. "FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age" *arXiv preprint arXiv:1908.04913*, 2019. https://arxiv.org/abs/1908.04913.
8. Varma, Sudhir and Richard Simon. “Bias in error estimation when using cross-validation for model selection.” *BMC Bioinformatics*, vol. 7, no. 1, 2006. https://doi.org/10.1186/1471-2105-7-91.
9. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.
10. Bland, J. Martin, and Douglas G. Altman. "Statistical Methods for Assessing Agreement between Two Methods of Clinical Measurement." *The Lancet*, vol. 327, no. 8476, 1986, pp. 307-310. https://doi.org/10.1016/S0140-6736(86)90837-8.
11. Pearson, Karl. (1901). LIII. On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science. https://doi.org/10.1080/14786440109462720.
12. Belhumeur, P.N.; Hespanha, J.P.; Kriegman, D.J.. (1997). Eigenfaces vs. Fisherfaces: recognition using class specific linear projection. IEEE Transactions on Pattern Analysis and Machine Intelligence. https://doi.org/10.1109/34.598228.
13. Higham, Nicholas J.. (2002). Accuracy and Stability of Numerical Algorithms: Second Edition. Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9780898718027.
14. Harris, Charles R.; Millman, K. Jarrod; van der Walt, Stéfan J.; Gommers, Ralf; Virtanen, Pauli; Cournapeau, David; Wieser, Eric; Taylor, Julian; Berg, Sebastian; Smith, Nathaniel J.; Kern, Robert; Picus, Matti; Hoyer, Stephan; van Kerkwijk, Marten H.; Brett, Matthew; Haldane, Allan; del Río, Jaime Fernández; Wiebe, Mark; Peterson, Pearu; Gérard-Marchant, Pierre; Sheppard, Kevin; Reddy, Tyler; Weckesser, Warren; Abbasi, Hameer; Gohlke, Christoph; Oliphant, Travis E.. (2020). Array programming with NumPy. Nature. https://doi.org/10.1038/s41586-020-2649-2.
15. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.
16. Stevens, S. S.. (1946). On the Theory of Scales of Measurement. Science. https://doi.org/10.1126/science.103.2684.677.
17. Kim, Minchul, et al. "AdaFace: Quality Adaptive Margin for Face Recognition." *arXiv preprint arXiv:2204.00964*, 2022. https://arxiv.org/abs/2204.00964.
18. Kemelmacher-Shlizerman, Ira, et al. "The MegaFace Benchmark: 1 Million Faces for Recognition at Scale" *arXiv preprint arXiv:1512.00596*, 2015. https://arxiv.org/abs/1512.00596.
19. Lin, Lawrence I-Kuei. (1989). A Concordance Correlation Coefficient to Evaluate Reproducibility. Biometrics. https://doi.org/10.2307/2532051.

20. Kapoor, Sayash; Narayanan, Arvind. (2023). Leakage and the reproducibility crisis in machine-learning-based science. Patterns. https://doi.org/10.1016/j.patter.2023.100804.

21. Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Millman, K. Jarrod; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, C J; Polat, İlhan; Feng, Yu; Moore, Eric W.; VanderPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R.; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; SciPy 1.0 Contributors; Vijaykumar, Aditya; Bardelli, Alessandro Pietro; Rothberg, Alex; Hilboll, Andreas; Kloeckner, Andreas; Scopatz, Anthony; Lee, Antony; Rokem, Ariel; Woods, C. Nathan; Fulton, Chad; Masson, Charles; Häggström, Christian; Fitzgerald, Clark; Nicholson, David A.; Hagen, David R.; Pasechnik, Dmitrii V.; Olivetti, Emanuele; Martin, Eric; Wieser, Eric; Silva, Fabrice; Lenders, Felix; Wilhelm, Florian; Young, G.; Price, Gavin A.; Ingold, Gert-Ludwig; Allen, Gregory E.; Lee, Gregory R.; Audren, Hervé; Probst, Irvin; Dietrich, Jörg P.; Silterra, Jacob; Webber, James T; Slavič, Janko; Nothman, Joel; Buchner, Johannes; Kulick, Johannes; Schönberger, Johannes L.; de Miranda Cardoso, José Vinícius; Reimer, Joscha; Harrington, Joseph; Rodríguez, Juan Luis Cano; Nunez-Iglesias, Juan; Kuczynski, Justin; Tritz, Kevin; Thoma, Martin; Newville, Matthew; Kümmerer, Matthias; Bolingbroke, Maximilian; Tartre, Michael; Pak, Mikhail; Smith, Nathaniel J.; Nowaczyk, Nikolai; Shebanov, Nikolay; Pavlyk, Oleksandr; Brodtkorb, Per A.; Lee, Perry; McGibbon, Robert T.; Feldbauer, Roman; Lewis, Sam; Tygier, Sam; Sievert, Scott; Vigna, Sebastiano; Peterson, Stefan; More, Surhud; Pudlik, Tadeusz; Oshima, Takuya; Pingel, Thomas J.; Robitaille, Thomas P.; Spura, Thomas; Jones, Thouis R.; Cera, Tim; Leslie, Tim; Zito, Tiziano; Krauss, Tom; Upadhyay, Utkarsh; Halchenko, Yaroslav O.; Vázquez-Baeza, Yoshiki. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. https://doi.org/10.1038/s41592-019-0686-2.

22. Meng, Qiang, et al. "MagFace: A Universal Representation for Face Recognition and Quality Assessment." *arXiv preprint arXiv:2103.06627*, 2021. https://arxiv.org/abs/2103.06627.

23. Huang, Gary B., Marwan Mattar, Tamara Berg, and Erik Learned-Miller. "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments." *ECCV Workshop on Faces in Real-Life Images*, 2008. https://cs.umass.edu/~elm/papers/Huang_eccv2008-lfw.pdf.

24. Wilkinson, Mark D.; Dumontier, Michel; Aalbersberg, IJsbrand Jan; Appleton, Gabrielle; Axton, Myles; Baak, Arie; Blomberg, Niklas; Boiten, Jan-Willem; da Silva Santos, Luiz Bonino; Bourne, Philip E.; Bouwman, Jildau; Brookes, Anthony J.; Clark, Tim; Crosas, Mercè; Dillo, Ingrid; Dumon, Olivier; Edmunds, Scott; Evelo, Chris T.; Finkers, Richard; Gonzalez-Beltran, Alejandra; Gray, Alasdair J.G.; Groth, Paul; Goble, Carole; Grethe, Jeffrey S.; Heringa, Jaap; 't Hoen, Peter A.C; Hooft, Rob; Kuhn, Tobias; Kok, Ruben; Kok, Joost; Lusher, Scott J.; Martone, Maryann E.; Mons, Albert; Packer, Abel L.; Persson, Bengt; Rocca-Serra, Philippe; Roos, Marco; van Schaik, Rene; Sansone, Susanna-Assunta; Schultes, Erik; Sengstag, Thierry; Slater, Ted; Strawn, George; Swertz, Morris A.; Thompson, Mark; van der Lei, Johan; van Mulligen, Erik; Velterop, Jan; Waagmeester, Andra; Wittenburg, Peter; Wolstencroft, Katherine; Zhao, Jun; Mons, Barend. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data. https://doi.org/10.1038/sdata.2016.18.

25. Cameron, A. Colin, and Douglas L. Miller. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. https://doi.org/10.3368/jhr.50.2.317.

26. Cao, Qiong, et al. "VGGFace2: A dataset for recognising faces across pose and age" *arXiv preprint arXiv:1710.08092*, 2017. https://arxiv.org/abs/1710.08092.

27. Meng, Zijie; Liu, Jiwen; Liu, Yufei; Tong, Chengzhuo; Liu, Xiaoqiang; Zhang, Yuanxing; Xu, Yulong; Wan, Pengfei. (2026). ARGUS: Stacked Multi-View Identity Mosaic Injection for Subject-Preserving Video Generation. arXiv:2606.11670. https://arxiv.org/abs/2606.11670.

28. Tipping, Michael E.; Bishop, Christopher M.. (1999). Probabilistic Principal Component Analysis. Journal of the Royal Statistical Society Series B: Statistical Methodology. https://doi.org/10.1111/1467-9868.00196.

29. AT&T Laboratories Cambridge. (2026). Modified Olivetti faces: Sam Roweis MATLAB version [Data archive; mirror accessed 7 November 2026]. File identity and row-label layout verified against the scikit-learn maintainer loader. https://ndownloader.figshare.com/files/5976027.

30. Turk, M.A.; Pentland, A.P.. (1991). Face recognition using eigenfaces. Proceedings. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr.1991.139758.