

405 **A Proof of Theorem 1**

406 We reiterate the setup and notation introduced in the paper here for ease of reference.

407 **Notation** $[N]$ denotes the set of natural number $\{1, 2, \dots, N\}$. Id denotes the (vector-valued)
 408 identity function. We write two functions f, g agreeing for all points in set P as $f \equiv_P g$. Finally,
 409 we write the total derivative of a vector-valued function f by all its inputs z as $\frac{\partial f}{\partial z}$, *i.e.* the Jacobian
 410 matrix with entries $\frac{\partial f_i}{\partial z_j}$.

411 **Setup** We are given two arbitrary distributions P, Q over latents $z = (z_1, \dots, z_K) \in \mathcal{Z}$. Each
 412 latent z_k describes one of the K *components* of the final data point x produced by the ground-truth
 413 data-generating process f . A model \hat{f} is trained to fit the data-generating process on samples of P ;
 414 the aim is to derive conditions on P and \hat{f} that are sufficient for \hat{f} to then also fit f on Q .

415 We assume that f, \hat{f} are chosen such that we can find at least one *compositional representations*
 416 (Definition 1) for either function that shares a common *composition function* C and factorization of
 417 the latent space $\mathcal{Z}_1 \times \dots \times \mathcal{Z}_K = \mathcal{Z}$.

418 *Proof of Theorem 1.* For \hat{f} to generalize to Q , we need to show fitting f on P implies also fitting it
 419 on Q , in other words

$$f \equiv_P \hat{f} \implies f \equiv_Q \hat{f} \quad (8)$$

420 *Step 1.* Since C is the same for both functions, we immediately get

$$\varphi \equiv_Q \hat{\varphi} \implies f \equiv_Q \hat{f}, \quad (9)$$

421 *i.e.* it suffices to show that the *component functions* generalize. Note, however, that since C is not
 422 generally assumed to be invertible, we do *not* directly get that agreement of f, \hat{f} on P also implies
 423 agreement of their component functions $\varphi, \hat{\varphi}$ on P .

424 *Step 2.* We require P to have *compositional support* w.r.t. Q (Definition 2 and Assumption (A2)).
 425 The consequence of this assumption is that any point $q = (q_1, \dots, q_K) \in Q$ can be constructed from
 426 components of the K *support points* $p^k = (p_1^k, \dots, p_K^k) \in P$ subject to $p_k^k = q_k$ as

$$q = (p_1^1, \dots, p_K^K). \quad (10)$$

427 A trivial consequence, then, is that points $\tilde{x} \in \tilde{\mathcal{X}}$ in *component space* corresponding to points in Q in
 428 latent space can always be mapped back to latents in P

$$\varphi(q) = (\varphi_1(q_1), \dots, \varphi_K(q_K)) = \left(\varphi_1(p_1^{(1)}), \dots, \varphi_K(p_K^{(K)}) \right) \quad (11)$$

429 because each *component function* φ_k only depends on the latents z_k of a single component. This is
 430 also the case for the component functions $\hat{\varphi}$ of \hat{f} so that we get

$$\varphi \equiv_P \hat{\varphi} \implies \varphi \equiv_Q \hat{\varphi}. \quad (12)$$

431 *Step 3.* We now only need to show that $\varphi \equiv_P \hat{\varphi}$ follows from $f \equiv_P \hat{f}$. As noted above, this is
 432 not guaranteed to be the case, as C is not generally invertible (*e.g.* in the presence of occlusions).
 433 We, therefore, need to consider when a unique reconstruction of the component functions φ (and
 434 correspondingly $\hat{\varphi}$) is possible, based on only the observations $x = f(z)$ on Q .

435 As explained in the main paper, we can reason about how a change in the latents z_k of some slot
 436 affects the final output, which we can express through the chain rule as

$$\underbrace{\frac{\partial f}{\partial z_k}}_{N \times D}(z) = \underbrace{\frac{\partial C}{\partial \varphi_k}}_{N \times M}(\varphi(z)) \underbrace{\frac{\partial \varphi_k}{\partial z_k}}_{M \times D}(z_k) \quad \forall k \in [K]. \quad (13)$$

437 Here, N is the dimension of the final output (e.g. $64 \times 64 \times 3$ for RGB images), M is the dimension
 438 of a component’s representation $\tilde{\mathbf{x}}_k$ (e.g. also $64 \times 64 \times 3$ for RGB images), and D is the dimension
 439 of a component’s latent description \mathbf{z}_k (e.g. 5: x-position, y-position, shape, size, hue for sprites).
 440 Note that we can look at the derivative component-wise because each *component function* φ_k only
 441 depends on the latents \mathbf{z}_k of its component. However, the *combination function* still depends on the
 442 (hidden) representation of all components, and therefore $\frac{\partial \mathbf{C}}{\partial \varphi_k}$ is a function of all φ and the entire \mathbf{z} .

443 In equation 13, the left-hand side (LHS) $\frac{\partial \mathbf{f}}{\partial \mathbf{z}_k}$ can be computed from the training, as long as $\text{supp } P$
 444 is an open set. On the right-hand side (RHS), the functional form of $\frac{\partial \mathbf{C}}{\partial \varphi_k}$ is known since \mathbf{C} is
 445 given, but since $\varphi(\mathbf{z})$ is still unknown, the exact entries of this Jacobian matrix are unknown. As
 446 such, equation 13 defines a system of partial differential equations (PDEs) for the set of component
 447 functions φ with independent variables \mathbf{z} .

448 Before we can attempt to solve this system of PDEs, we simplify it by isolating $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}$. Since all terms
 449 are matrices, this is equivalent to solving a system of linear equations. For $N = M$, $\frac{\partial \mathbf{C}}{\partial \varphi_k}$ is square,
 450 and we can solve by taking its inverse as long as the determinant is not zero. In the general case of
 451 $N \geq M$, however, we have to resort to the pseudoinverse to write

$$\frac{\partial \varphi_k}{\partial \mathbf{z}_k}^* = \left(\frac{\partial \mathbf{C}}{\partial \varphi_k}^\top \frac{\partial \mathbf{C}}{\partial \varphi_k} \right)^{-1} \frac{\partial \mathbf{C}}{\partial \varphi_k}^\top \frac{\partial \mathbf{f}}{\partial \mathbf{z}_k} \quad \forall k \in [K], \quad (14)$$

452 which gives all solutions $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}^*$ if any exist. This system is overdetermined, and a (unique) solution
 453 exists if $\frac{\partial \mathbf{C}}{\partial \varphi_k}$ has full (column) rank. In other words, to execute this simplification step on P , we
 454 require that for all $\mathbf{z} \in P$ the M column vectors of the form

$$\left(\frac{\partial C_1}{\partial \varphi_{km}}(\varphi(\mathbf{z})), \dots, \frac{\partial C_N}{\partial \varphi_{km}}(\varphi(\mathbf{z})) \right)^\top \quad \forall m \in [M] \quad (15)$$

455 are linearly independent. Each entry of a column vector describes how all entries C_n of the final
 456 output (e.g. the pixels of the output image) change with a single entry φ_{km} of the intermediate
 457 representation of component k (e.g. a single pixel of the component-wise image). It is easy to see
 458 that if even a part of the intermediate representation is not reflected in the final output (e.g. in the
 459 presence of occlusions, when a single pixel of one component is occluded), the entire corresponding
 460 column is zero, and the matrix does not have full rank.

461 To circumvent this issue, we realize that the LHS of equation 14 only depends on the latents \mathbf{z}_k of a
 462 single component. Hence, for a given latent \mathbf{z} and a slot index k , the correct component function will
 463 have the same solution for all points in the set

$$P'(\mathbf{z}, k) = \{\mathbf{p} \in \text{supp } P \mid \mathbf{p}_k = \mathbf{z}_k\}. \quad (16)$$

464 We can interpret these points as the intersection of P with a plane in latent space at \mathbf{z}_k (e.g. all latent
 465 combinations in the training set in which one component is fixed in a specific configuration). We can
 466 then define a modified composition function $\tilde{\mathbf{C}}$ that takes \mathbf{z} and a slot index k as input and produces
 467 a “superposition” of images corresponding to the latents in the subset as

$$\tilde{\mathbf{C}}(\mathbf{z}, k) = \sum_{\mathbf{p} \in P'(\mathbf{z}, k)} \mathbf{C}(\varphi(\mathbf{p})). \quad (17)$$

468 Essentially, we are condensing the information from multiple points in the latent space into a single
 469 function. This enables us to write a modified version of equation 13 as

$$\sum_{\mathbf{p} \in P'(\mathbf{z}, k)} \frac{\partial \mathbf{f}}{\partial \mathbf{z}_k}(\mathbf{p}) = \sum_{\mathbf{p} \in P'(\mathbf{z}, k)} \frac{\partial \mathbf{C}}{\partial \varphi_k}(\varphi(\mathbf{p})) \frac{\partial \varphi_k}{\partial \mathbf{z}_k}(\mathbf{z}_k) = \frac{\partial \tilde{\mathbf{C}}}{\partial \varphi_k}(\mathbf{z}, k) \frac{\partial \varphi_k}{\partial \mathbf{z}_k}(\mathbf{z}_k) \quad \forall k \in [K] \quad (18)$$

470 Now we can solve for $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}$ as in equation 14, but this time require only that $\frac{\partial \tilde{\mathbf{C}}}{\partial \varphi_k}$ has full (column)
 471 rank for a unique solution to exist, i.e.

$$\text{rank} \frac{\partial \tilde{\mathbf{C}}}{\partial \varphi_k}(\mathbf{z}, k) = \sum_{\mathbf{p} \in P'(\mathbf{z}, k)} \frac{\partial \mathbf{C}}{\partial \varphi_k}(\varphi(\mathbf{p})) = M \quad \forall \mathbf{z} \in P \quad \forall k \in [K]. \quad (19)$$

472 In general, this condition is easier to fulfill since full rank is not required in any one point but over
 473 a set of points. For occlusions, for example, any pixel of one slot can be occluded in some points
 474 $\mathbf{p} \in P'$, as long as it is not occluded in all of them. We can interpret this procedure as “collecting
 475 sufficient information” such that an inversion of the generally non-invertible C becomes feasible
 476 locally.

477 The requirement that $\text{supp } P$ has to be an open set, together with the full rank condition on the
 478 Jacobian of the composition function condensed over multiple points, \tilde{C} , is termed *sufficient support*
 479 in the main paper (Definition 3 and Assumption (A3)). As explained here, this allows for the
 480 reconstruction of $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}$ from the observations, *i.e.*

$$\mathbf{f} \stackrel{\equiv}{P} \hat{\mathbf{f}} \implies \frac{\partial \varphi}{\partial \mathbf{z}} \stackrel{\equiv}{P} \frac{\partial \tilde{\varphi}}{\partial \mathbf{z}}. \quad (20)$$

481 *Step 4.* The above step only gives us agreement of the *derivative* of the component functions, $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}$,
 482 not agreement of the functions themselves. As explained above, the solution to the linear system of
 483 equations 14 constitutes a system of partial differential equations (PDEs) in the set of component
 484 functions φ with independent variables \mathbf{z} . We can see that this system has the form

$$\partial_i \varphi(\mathbf{z}) = \mathbf{a}_i(\mathbf{z}, \varphi(\mathbf{z})), \quad (21)$$

485 where $i \in [L] = [K \times D]$ is an index over the flattened dimensions K and D such that $\partial_i \varphi$ denotes
 486 $\frac{\partial \varphi}{\partial z_L}$ (which is essentially one column of $\frac{\partial \varphi_k}{\partial \mathbf{z}_k}$ aggregated over all k) and \mathbf{a}_i is the combination of
 487 corresponding terms from the LHS. If this system allows for more than one solution, we cannot
 488 uniquely reconstruct the component functions from their derivatives.

489 If we have access to some initial point, however, for which we know $\varphi(\mathbf{0}) = \varphi^0$, we can write

$$\begin{aligned} \varphi(z_1, \dots, z_L) - \varphi^* &= (\varphi(z_1, \dots, z_L) - \varphi(0, z_2, \dots, z_L)) \\ &\quad + (\varphi(0, z_2, \dots, z_L) - \varphi(0, 0, z_3, \dots, z_L)) \\ &\quad + \dots \\ &\quad + (\varphi(0, \dots, 0, z_L) - \varphi(0, \dots, 0)). \end{aligned} \quad (22)$$

490 In each line of this equation, only a single $z_i =: t$ is changing; all other z_1, \dots, z_L are fixed. Any
 491 solution of 22, therefore, also has to solve the L ordinary differential equations (ODEs) of the form

$$\partial_i \varphi(z_1, \dots, z_{i-1}, t, z_{i+1}, \dots, z_L) = \mathbf{a}_i(z_1, \dots, z_{i-1}, t, z_{i+1}, \dots, z_L, \varphi(z_1, \dots, z_{i-1}, t, z_{i+1}, \dots, z_L)), \quad (23)$$

492 which have a unique solution if \mathbf{a}_i is Lipschitz in φ and continuous in z_i , as guaranteed by (A1).
 493 Therefore, 22 has at most one solution. This reference point does not have to be in $\mathbf{z} = \mathbf{0}$, as a simple
 494 coordinate transform will yield the same result for any point in P . It is therefore sufficient that there
 495 exists *some* point $\mathbf{p}^0 \in P$ for which $\varphi(\mathbf{p}^0) = \hat{\varphi}(\mathbf{p}^0)$ to obtain the same unique solution for φ and
 496 $\hat{\varphi}$, which is exactly what (A4) states. Overall, this means that agreement of the derivatives of the
 497 component functions also implies agreement of the component functions themselves, *i.e.*

$$\frac{\partial \varphi}{\partial \mathbf{z}} \stackrel{\equiv}{P} \frac{\partial \tilde{\varphi}}{\partial \mathbf{z}} \implies \varphi \stackrel{\equiv}{P} \hat{\varphi} \quad (24)$$

498 *Step 5.* Finally, we can conclude the model $\hat{\mathbf{f}}$ fitting the ground-truth generating process \mathbf{f} on the
 499 training distribution P , through 20, 24, 12, 9, implies the model generalizing to Q as well. In other
 500 words, equation 8 holds.

501 □

502 B Details about the compositional functions

503 As explained in equation 7 in section 4, the composition function is implemented as a soft pixel-wise
 504 addition in most experiments. The use of the sigmoid function $\sigma(\cdot)$ in the composition

$$\mathbf{x} = \sigma(\tilde{\mathbf{x}}_1) \cdot \tilde{\mathbf{x}}_1 + \sigma(-\tilde{\mathbf{x}}_1) \cdot \tilde{\mathbf{x}}_2 \quad (25)$$

505 was necessary for training stability. With this formulation, sprites can also overlap somewhat
 506 transparently, which is not desired and leads to small reconstruction artifacts for some specific
 507 samples. Implementing the composition with a step function as

$$\mathbf{x} = \text{step}(\tilde{\mathbf{x}}_1) \cdot \tilde{\mathbf{x}}_1 + \text{step}(-\tilde{\mathbf{x}}_1) \cdot \tilde{\mathbf{x}}_2 \quad (26)$$

508 instead would be more faithful to the ground-truth data-generating process, but is hard to train with
 509 gradient descent.

510 Note that both formulations could easily be extended to more than one sprite by simply repeating the
 511 composition operation with any additional sprite.

512 In section 4, we also looked at a model that implements the composition through alpha compositing
 513 instead (see also Table 1, #11). Here, each component’s intermediate representation is an RGBA
 514 image. The components are then overlaid on an opaque black background using the composition
 515 function

$$x_\alpha = x_{1,\alpha} + (1 - x_{1,\alpha}) \cdot x_{2,\alpha} \quad (27)$$

$$x_{\text{RGB}} = x_{1,\alpha} \cdot x_{1,\text{RGB}} + (1 - x_{1,\alpha}) \cdot \frac{x_{2,\alpha}}{x_\alpha} \cdot x_{2,\text{RGB}}. \quad (28)$$

516 While this yields a compositional function, the sufficient support condition (Definition 3) is generally
 517 not fulfilled on the sprites data. The reason is that in fully transparent pixels ($\alpha = 0$), changing the
 518 RGB value is not reflected in the output. Conversely, if a pixel is black, changing its alpha value
 519 will not affect how it is blended over a black background. As a result, most columns in the Jacobian
 520 $\frac{\partial \mathcal{C}}{\partial \varphi_k}$ (see also equation 15) will be zero. Since the intermediate representations of each sprite will
 521 contain a lot of black or transparent pixels (the entire background), the rank of the Jacobian here will
 522 be low. In this case, the workaround from equation 17 does not help since the low rank is not a result
 523 of another component in the foreground but of the specific parameterization of each component itself.

524 As stated in the main paper, the fact that this parameterization still produces good results and
 525 generalizes well is an indicator that there might be another proof strategy or workaround that avoids
 526 this specific issue.