

Compressed Diffusion

Scott Gigante^{1 †}, Jay S. Stanley III^{1 †}, Ngan Vu², David van Dijk³,
Kevin R. Moon⁴, Guy Wolf^{5 †† §§}, and Smita Krishnaswamy^{2 3 †† §}

¹ Comp. Bio. & Bioinformatics Program, ² Dept. of Comp. Sci., ³ Dept. of Genetics, Yale University, New Haven, CT, USA; ⁴ Dept. of Mathematics and Statistics, Utah State University, Logan, UT, USA
⁵ Dept. of Mathematics and Statistics, Université de Montréal, Montréal, QC, Canada

Abstract—Diffusion maps are a commonly used kernel-based method for manifold learning, which can reveal intrinsic structures in data and embed them in low dimensions. However, as with most kernel methods, its implementation requires a heavy computational load, reaching up to cubic complexity in the number of data points. This limits its usability in modern data analysis. Here, we present a new approach to computing the diffusion geometry, and related embeddings, from a compressed diffusion process between data regions rather than data points. Our construction is based on an adaptation of the previously proposed measure-based Gaussian correlation (MGC) kernel that robustly captures the local geometry around data points. We use this MGC kernel to efficiently compress diffusion relations from pointwise to data region resolution. Finally, a spectral embedding of the data regions provides coordinates that are used to interpolate and approximate the pointwise diffusion map embedding of data. We analyze theoretical connections between our construction and the original diffusion geometry of diffusion maps, and demonstrate the utility of our method in analyzing big datasets, where it outperforms competing approaches.

I. INTRODUCTION

Manifold learning approaches are often used for modeling and uncovering intrinsic low dimensional structure in high dimensional data (e.g., in genomics [7]). Diffusion maps [3], in particular, are a popular method that capture data manifolds with random walks that propagate through nonlinear pathways in the data. Transition probabilities of a Markovian diffusion process define an intrinsic diffusion distance metric that is amenable to low dimensional embedding. Indeed, by arranging transition probabilities in a row-stochastic diffusion operator, and taking its leading eigenvalues and eigenvectors, one can derive a small set of coordinates where diffusion distances are approximated as Euclidean distances, and intrinsic manifold structures are revealed.

While the embedding provided by diffusion maps is useful for data analysis, it is also challenging to apply to modern “big data” settings due to scalability issues. As is often the case with kernel methods, diffusion maps require the computation of pointwise transition probabilities, which is computationally quadratic in the size of the data, and an eigendecomposition which is cubic if all eigenvectors are computed. Here, we show that one can reduce the computational cost significantly by

only considering transition probabilities between data regions, which can be efficiently computed from a compressed affinity kernel over a fixed-size partition of the data. Our construction is based on coarse-graining a measure-based kernel [1], [2] with an inverse-density measure, which was recently introduced in [5] for geometry-based data generation and uniform resampling of data manifolds. Here we extend the uses of this kernel to provide efficient implementation and application of diffusion maps by first embedding data regions, and then interpolating the pointwise embedding from their diffusion coordinates.

The main contributions of this work are as follows. On the theoretical side, we further establish the relations between the original diffusion framework from [3] and the construction in [5], both at a pointwise and compressed (i.e., data-region) level. On the practical side, we suggest a novel partitioning method, the results of which indicate significant speedups in the computation of the diffusion embedding, which outperform other approaches, and enables the application of diffusion-based manifold learning well beyond the data sizes traditionally used with kernel methods.

II. PROBLEM SETUP

A. Preliminaries

Let $\mathcal{M} \subseteq \mathbb{R}^m$ be a compact d dimensional manifold immersed in the ambient space \mathbb{R}^m , where $d \ll m$, which represents the intrinsic geometry of data sampled from it. For simplicity, the integration notation $\int \cdot dy$ in this paper will refer to the Lebesgue integral $\int_{\mathcal{M}} \cdot dy$ over the manifold, instead of the whole space \mathbb{R}^n . Further, while (for simplicity) such integrals are written without a specific measure one can equivalently, w.l.o.g., replace dx with an appropriate measure representing data sampling distribution over \mathcal{M} . Let $g(x, y) \triangleq \exp(-\|x - y\|^2/\varepsilon)$, $x, y \in \mathcal{M}$, $\varepsilon > 0$, define the Gaussian kernel $Gf(x) = \int g(x, y)f(y)dy$ used in [3] to capture local neighborhoods from data sampled from \mathcal{M} . Following [3] and related work, we define the Gaussian degree $q(x) = \|g(x, \cdot)\|_1 = \int g(x, y)dy$ and assume it provides a suitable approximation of the distribution (or local density) of data over the manifold \mathcal{M} . Finally, given a measure μ over the manifold, an MGC kernel [1], [2] is defined as $k_\mu(x, y) = \int g(x, r)g(y, r)d\mu(r)$. Note that while we use a Gaussian kernel for the remainder of this work, the definitions and theorems to follow do not depend on the choice of g , so long as it is a kernel function.

[†] These authors contributed equally. ^{††} These authors contributed equally.

[§] Corresponding author. smita.krishnaswamy@yale.edu
333 Cedar St. New Haven CT 06510 USA

^{§§} Corresponding author. guy.wolf@umontreal.ca
Pav. André-Aisenstadt, 2920 ch.d.l. Tour, Montréal, QC H3T 1J4, Canada

The original MGC construction in [2] considered measures that represent data distribution, and used the constructed MGC kernel to define diffusion maps (see Sec. II-B and [3]) with them. Recently, it was shown in [5] that an MGC kernel constructed with inverse-density measure allows separation of data geometry and distribution, which in turn allows uniform data generation over the data manifold, with applications in alleviating sampling biases in data analysis (e.g., imbalanced classification). Here, we further explore the inverse-density MGC kernel and its properties. In particular, we show this kernel enables compression of its resulting diffusion geometry into data regions, instead of data points, to efficiently capture the intrinsic manifold geometry of analyzed data.

B. Diffusion Maps

Diffusion maps [3] utilize a set of local affinities to define a Markovian diffusion process over analyzed data, which captures the intrinsic data geometry via diffusion distances. The original construction [3] defines transition probabilities between data points based on Gaussian affinities as $p(x,y) = g(x,y)/q(x)$, where it can be verified that $\int p(x,y)dy = 1$. This construction can also be generalized to other affinity kernels, such as the MGC kernel $k_\mu(x,y)$ from [2] or variation in [5]. Under mild conditions on the affinity kernel, the resulting transition probability operator has a discrete decaying spectrum of eigenvalues $1 = \lambda_0 \geq |\lambda_1| \geq |\lambda_2| \geq \dots$ which are used together with their corresponding eigenfunctions $\phi_0, \phi_1, \phi_2, \dots$ (with ϕ_0 being constant) to achieve the diffusion map of the data. Each data point $x \in \mathcal{M}$ is embedded by this diffusion map to the diffusion coordinates $\Phi_t(x) = (\lambda_1^t \phi_1(x), \dots, \lambda_\delta^t(x) \phi_\delta(x))$, where the exact value of δ depends on the spectrum of the transition probabilities operator $Pf(x) = \int p(x,y)f(y)dy$, whose kernel is $p(x,y)$.

III. INVERSE-DENSITY MGC KERNEL

Our construction here is based on an inverse-density MGC kernel, first introduced in [5], which is defined as follows:

Definition 1. *The inverse-density MGC (ID-MGC) kernel is defined as $k(x,y) = \int \frac{g(x,r)g(r,y)}{q(r)} dr$, $x,y \in \mathcal{M}$, with the associated integral operator $Kf(x) = \int k(x,y)f(y)dy$.*

This kernel corresponds to $k_\mu(x,y)$ with $d\mu(x) = q^{-1}(x)dx$, where in practice $q(x)$ accounts (up to normalization) for data density over \mathcal{M} . Therefore, the connectivity captured by this kernel normalizes density variations by enhancing relations in sparse regions compared to dense ones. Indeed, in [5] this property was used to adjust the distribution of data and provide uniform resampling of data manifolds. We note that since we consider \mathcal{M} as representing the intrinsic geometry of collected (or observed) data, we expect a certain amount of data points to exist in each local region, and thus $q^{-1}(x)$ can be expected to have a finite upper bound over \mathcal{M} .

Several useful properties of the ID-MGC kernel and its relation to the Gaussian-based diffusion operator in Sec. II-B are summarized in the following theorem (proof omitted for brevity, but can be seen in [4]):

Theorem 1. *Let the kernel $k(\cdot, \cdot)$ and operator K be defined as in Def. 1. Then, this kernel (and operator) can be related to the Gaussian kernel G and diffusion operator P via 1) the operator norm: $\|K\| \leq \|G\|$, 2) the kernel degrees: $\|k(x, \cdot)\|_1 = \|g(x, \cdot)\|_1$, $x \in \mathcal{M}$, and 3) two-step diffusion: $P^2 f(x) = \int \frac{k(x,y)}{\|k(x, \cdot)\|_1} f(y)dy$.*

IV. COMPRESSED KERNEL AND DIFFUSION TRANSITIONS

Let $S_1, \dots, S_n \subset \mathcal{M}$ be a partition of the manifold into measurable subsets, such that $\mathcal{M} = \bigcup_{j=1}^n S_j$, and $S_i \cap S_j = \emptyset$ for every $i \neq j$. We define a compressed kernel over such a partition as follows:

Definition 2. *Let $k(x,y)$, $x,y \in \mathcal{M}$ be defined as in Def. 1. The compressed kernel over partition $\mathcal{S} = \{S_j\}_{j=1}^n$ of \mathcal{M} is given by*

$$k_{\mathcal{S}}(S,T) = \int_S \int_T k(x,y) dx dy, \quad S,T \in \mathcal{S},$$

with the corresponding $n \times n$ kernel matrix given by

$$[K_{\mathcal{S}}]_{ij} = k(S_i, S_j), \quad i,j = 1, \dots, n.$$

Similar to the original diffusion map construction in Sec. II-B, we normalize the compressed kernel to get diffusion probabilities $p_{\mathcal{S}}(S,T) = k(S,T)/\|k(S, \cdot)\|_1$, organized in an $n \times n$ row-stochastic matrix $P_{\mathcal{S}}$. This matrix captures diffusion transition probabilities between data regions. The relation between the compressed construction and the original diffusion framework is summarized in the following theorem (proof omitted for brevity, but can be seen in [4]):

Theorem 2. *The compressed construction here can be related to the pointwise one in Sec. II-B, via 1) operator norm (in matrix & operator form): $\|K_{\mathcal{S}}\| \leq O(\|G\|)$ with a constant that only depends on the finite volume of \mathcal{M} 2) kernel degrees: $\|k_{\mathcal{S}}(S, \cdot)\|_1 = \int_S \|g(x, \cdot)\|_1 dx$, $S \in \mathcal{S}$, and 3) diffusion probabilities: $P_{\mathcal{S}}(S,T) = \int \Pr[S \overset{1 \text{ step}}{\rightsquigarrow} r] \Pr[r \overset{1 \text{ step}}{\rightsquigarrow} T] dr$ where $\Pr[r \overset{1 \text{ step}}{\rightsquigarrow} T] = \int_T p(r,y)dy$ and $\Pr[S \overset{1 \text{ step}}{\rightsquigarrow} r] = \int_S p(x,r) \Pr[x|S] dx$ with prior $\Pr[x|S] = \frac{\|g(x, \cdot)\|_1}{\|k_{\mathcal{S}}(S, \cdot)\|_1}$.*

V. COMPRESSION-BASED FAST DIFFUSION MAPS

The compressed construction of $K_{\mathcal{S}}$ and $P_{\mathcal{S}}$ gives rise to a natural approximation of diffusion maps. Given a partitioning \mathcal{S} , we define a compressed diffusion map $\Phi_{\mathcal{S}}^t : \mathcal{S} \rightarrow \mathbb{R}^d$ analogous to the pointwise one in II-B using eigenvectors of $P_{\mathcal{S}}$ (with corresponding eigenvalues) in lieu of the eigenfunctions of P . This diffusion map provides an embedding of data regions, rather than points, which represents a coarse version of the diffusion geometry over \mathcal{M} . Then, using the region-to-point transition probabilities $\Pr[S \overset{1 \text{ step}}{\rightsquigarrow} x]$ from Theorem 2, we approximate the pointwise diffusion map as $\tilde{\Phi}^t(x) = \sum_{j=1}^n \Phi_{\mathcal{S}}^t(S_j) \Pr[S_j \overset{1 \text{ step}}{\rightsquigarrow} x]$. Results in Sec. VI indicate that with the proper choice of partitions in \mathcal{S} , the compressed diffusion map $\tilde{\Phi}^t$ provides a good approximation of the original pointwise one from [3], while also providing

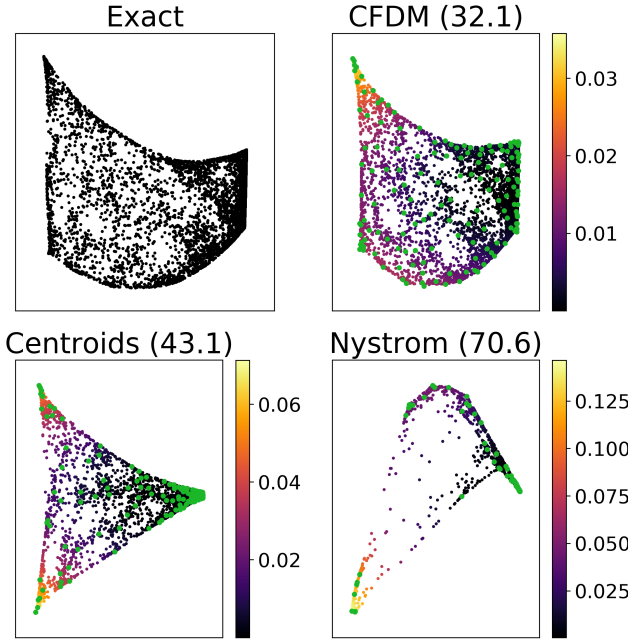


Fig. V.1. Example of approximations of diffusion maps on the Swiss Roll. Landmarks or centroids are shown in green. Mean squared error (MSE) of each approximation is shown in parentheses. Color shows MSE for each point.

significant scalability and performance advantages over both direct computation and other alternative approximations.

It now remains to derive a strategy for efficiently partitioning the data manifold such that compression over \mathcal{S} will suitably capture (albeit at a coarser level) the pointwise diffusion geometry defined by P . To this end, it is convenient to consider diffusion affinities and distances [1]–[3], rather than transition probabilities. These are defined as follows:

Definition 3. The diffusion affinity [3] kernel is the symmetric conjugate of P , which is defined spectrally as $Af(x) = \sum_{j=1}^{\infty} \lambda_j \langle \psi_j, f \rangle \psi_j(x)$ with $\psi_j = q^{1/2}(x) \phi_j(x)$.

Definition 4. The diffusion distance [3] between two points x and y sampled from \mathcal{M} is defined via the kernel function $a^t(\cdot, \cdot)$ of A^t , and equivalently via its eigendecomposition, as

$$\begin{aligned} [D^t(x, y)]^2 &= D_2^t(x, y) = \left\| a^t(x, \cdot) - a^t(y, \cdot) \right\|_2^2 \\ &= \sum_{j=0}^{\infty} \lambda_j^{2t} [\psi_j(x) - \psi_j(y)]^2 \end{aligned}$$

We note that in the described compressed diffusion scheme, we replace pointwise affinities with affinities between data points and partitions, rather than just between partitions. Under proper partitioning of the data, we expect (or assume) such substitution would retain intrinsic structure of the data, as captured by the diffusion geometry defined by P or A . The following proposition examines the difference between pointwise and coarse-grained diffusion affinity information, and relates this difference to locality of partitions with respect to the diffusion distance metric.

Proposition 3. Let $S \in \mathcal{S}$ and $\xi > 0$. If the diffusion affinities satisfy $\sup_{z \in \mathcal{M}, x \in S_j} |a(x, z) - \mathbb{E}_{y \in S_j}[a(y, z)]| < \xi$, then the

diffusion distances $D^1(x, y)$, $x, y \in S$, are bounded from above by $O(\xi)$, with a constant that only depends on $\text{vol } \mathcal{M}$.

Proof. For any $x, y \in S_j$ and $z \in \mathcal{M}$, due to the triangle inequality, we have $|a(x, z) - a(y, z)| \leq |a(x, z) - \mathbb{E}_{u \in S_j}[a(u, z)]| + |a(y, z) - \mathbb{E}_{u \in S_j}[a(u, z)]| \leq 2\xi$. Therefore, $\int |a(x, z) - a(y, z)|^2 dz \leq 4\xi^2 \text{vol } \mathcal{M}$, which together with Definition 4 yields the result in the proposition. \square

Indeed, this result indicates that to faithfully capture the diffusion geometry by given partitioning \mathcal{S} , the partitions should be local in the diffusion geometry, since the diffusion geometry in turn preserves the local geometry of the manifold [3]. This understanding motivates our formulation of a partitioning strategy, as explained below.

Here we turn to practical settings of data analysis applications. Therefore, in the following we consider finite settings, particularly for some dataset $X \subset \mathcal{M}$, $|X| < \infty$, sampled from the manifold, with P and A being $|X| \times |X|$ matrices constructed from it. Note also that in this setting, the eigenfunctions of P or A become eigenvectors (i.e., in $\mathbb{R}^{|X|}$), and the RHS in Definition 4 is written analogously with the sum going over $|X|$ eigenpairs.

In order to define a partition on X , we first choose a set of partition centroids L . Noting the success of coherence-based sampling strategies for signal compression and recovery [8], we propose a random sampling without replacement biased by coherence, where we estimate the coherence as follows.

Proposition 4. Let $\mu_\ell(x) = \sum_{j=0}^{\ell} |\psi_j(x)|^2$ be the order- ℓ coherence of $x \in X$ as in [8]. Then, the t -step diffusion coherence $\rho_t(x) = \|a^t(x, \cdot)\|_2^2$ is an estimate of $\mu_\ell(x)$ with error

$$\|\mu_\ell(x) - \rho_t(x)\|_2^2 \propto \left| \ell - \sum_{j=0}^{|X|-1} \lambda_j^{2t} \right|$$

for every $x \in X$.

The proof is omitted for brevity, but can be seen in [4].

According to Proposition 3, we should assign $x \in X$ to the partition associated with centroid $y \in L$ such that $D_2^t(x, y)$ is minimized. However, this distance is in practice biased by the diffusion coherence $\rho_t(y)$. We therefore define the *angular diffusion distance*, which we will show maximizes transition probabilities between points and the assigned centroids.

Definition 5. The angular diffusion distance between $x, y \in \mathcal{M}$ is defined using the eigendecomposition of the diffusion affinity kernel from Def. 3 as

$$D_c^t(x, y) = \arccos \left(\frac{\sum_{j=0}^{\infty} \lambda_j^{2t} \psi_j(x) \psi_j(y)}{\sqrt{\rho_t(x) \rho_t(y)}} \right).$$

Proposition 5. For $x \in \mathcal{M}$, the following dualities hold:

$$1) \arg \min_{y \in \mathcal{M}} D_c^t(x, y) = \arg \max_{y \in \mathcal{M}} \frac{a^{2t}(x, y)}{\sqrt{\rho_t(x) \rho_t(y)}};$$




Fig. VI.1. Quantitative evaluations. *Left*: Runtime scaling for 150 partitions over an increasing number of points. *Middle*: Approximation error on 32 diffusion components over an increasing number of partitions. *Right*: Approximation error compared to runtime over 2^{14} points.

$$2) \arg \min_{y \in \mathcal{M}} \left\| \frac{a^t(x, \cdot)}{\sqrt{\rho_t(x)}} - \frac{a^t(y, \cdot)}{\sqrt{\rho_t(y)}} \right\|_2^2 = \arg \min_{y \in \mathcal{M}} D_c^t(x, y).$$

The proof is omitted for brevity, but can be seen in [4]. Propositions 4 and 5 provide a convenient method to select a set of partitions that optimize proposition 3. To do this, we assign points $x \in X$ to partitions $S_i \in \mathcal{S}$ corresponding to centroids $y_i \in L$ according to

$$S_i = \{x \in X : D_c^t(x, y_i) = \arg \min_{y \in L} D_c^t(x, y)\}.$$

VI. EMPIRICAL RESULTS

A. Swiss roll

Here we compare three techniques for approximation of diffusion maps: linear interpolated diffusion on cluster centroids, volume-weighted Nystrom extension [6], and our proposed compression-based fast diffusion map (CFDM).

The Swiss roll is a canonical test dataset for diffusion maps, and manifold learning in general. Figure V.1 shows the diffusion map of the Swiss roll and each approximation, with the partition diffusion coordinates plotted in green. To quantify the reconstruction, we calculated the sum of squared error (SSE) across 32 diffusion components of 200 Swiss rolls, allowing for sign flips and reordering of components. CFDM produces a lower approximation error and a smaller approximate kernel for less computational cost than competing methods (Fig. VI.1). For $n \geq 10^4$, CFDM offers an approximately 10x speed-up over exact diffusion maps, and this advantage grows slightly as n increases.

B. Mass cytometry of induced pluripotent stem cells

Mass cytometry is the measurement of protein abundance in individual cells via mass spectrometry. We used CFDM to approximate the diffusion map of 20,000 single cells in an induced pluripotent stem cell (iPSC) system [9] with 500 partitions. Figure VI.2 shows the exact and approximated diffusion maps. CFDM produces a visually equivalent embedding in a fraction of the time, accurately revealing the differentiation of skin cells into both iPSCs and a failed reprogramming state. Such an approximation will be beneficial for constructing diffusion maps on large graphs, such as social networks, as well as for downstream applications, where algorithms of high computational complexity are run on the diffusion map.

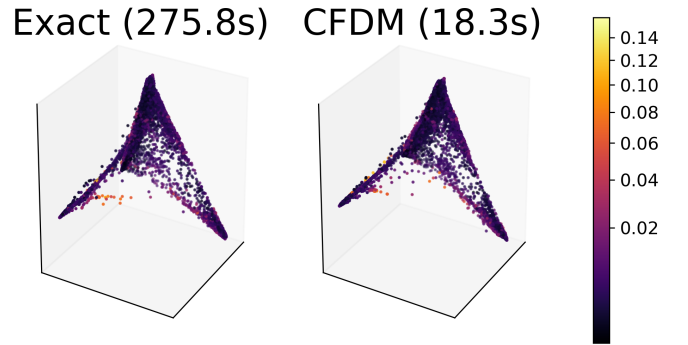


Fig. VI.2. Exact and approximate diffusion map of iPSC dataset. Runtime for 100 components shown above. Plots are colored by SSE of approximation.

Acknowledgments

This research was partially funded by: the Gruber Foundation [S.G.]; IVADO (l'institut de valorisation des données) [G.W.]; and the Chan-Zuckerberg Initiative (grant ID: 182702) [S.K.].

REFERENCES

- [1] A. Bermanis, G. Wolf, and A. Averbuch. Measure-based diffusion kernel methods. In *Proc. of the 10th intl. conf. on Sampling Theory and Applications (SampTA 2013)*, pages 489–492, Bremen, Germany, 2013.
- [2] A. Bermanis, G. Wolf, and A. Averbuch. Diffusion-based kernel methods on Euclidean metric measure spaces. *Appl. and Comp. Harm. Analysis*, 41(1):190–213, 2016.
- [3] R.R. Coifman and S. Lafon. Diffusion maps. *Appl. and Comp. Harm. Analysis*, 21(1):5–30, 2006.
- [4] S. Gigante, J.S. Stanley III, N. Vu, D. van Dijk, K.R. Moon, G. Wolf, and S. Krishnaswamy. http://krishnaswamylab.org/s/cfdm_proofs, 2018.
- [5] O. Lindenbaum, J.S. Stanley, G. Wolf, and S. Krishnaswamy. Geometry-based data generation. In *Adv. in Neural Info. Proc. Systems 31 (NeurIPS 2018)*, pages 1405–1416, Montréal, QC, Canada, 2018.
- [6] A.W Long and A.L Ferguson. Landmark diffusion maps (L-dMaps): Accelerated manifold learning out-of-sample extension. *Appl. and Comp. Harm. Analysis*, 47(1):190–211, 2019.
- [7] K.R. Moon, J.S. Stanley, D. Burkhardt, D. van Dijk, G. Wolf, and S. Krishnaswamy. Manifold learning-based methods for analyzing single-cell RNA-sequencing data. *Curr. Op. in Sys. Bio.*, 7:36–46, 2018.
- [8] G. Puy, N. Tremblay, R. Gribonval, and P. Vandergheynst. Random sampling of bandlimited signals on graphs. *Appl. and Comp. Harm. Anal.*, 44(2):446–475, 2018.
- [9] E.R Zunder, E. Lujan, Y. Goltsev, M. Wernig, and G.P. Nolan. A continuous molecular roadmap to iPSC reprogramming through progression analysis of single-cell mass cytometry. *Cell Stem Cell*, 16(3):323–337, 2015.