

ChangeFlow - Latent Rectified Flow for Change Detection in Remote Sensing

Blaž Rolih, Matic Fučka, Filip Wolf, and Luka Čehovin Zajc

University of Ljubljana, Faculty of Computer and Information Science, Slovenia
{ blaz.rolih@fri.uni-lj.si}

Abstract. Remote sensing change detection (RSCD) aims to localise changes between two images of the same geographic region. In practice, change masks often follow region-level annotation conventions rather than purely local appearance differences, making them context-dependent and occasionally ambiguous. Most state-of-the-art methods utilise per-pixel discriminative classification, which produces a single prediction per input and fails to explicitly model the changed region as a coherent whole. A natural alternative is generative formulation, which can model a distribution of plausible masks, enabling sampling to capture ambiguity and encourage global consistency. However, existing generative RSCD approaches typically lag behind strong discriminative baselines due to the high computational cost of pixel-space generation and the complexity of their conditioning mechanisms. To address the limitations of prior discriminative and generative methods, we propose **ChangeFlow**, a generative framework that reformulates change detection as the synthesis of a change mask in latent space via rectified flow. ChangeFlow is guided by a structured yet lightweight conditioning signal, and its stochastic design naturally supports sampling-based prediction ensembling. Namely, aggregating multiple predicted change masks improves robustness, while sample agreement provides a practical confidence estimation that highlights ambiguous regions. Across four benchmarks, ChangeFlow achieves an average F1 of 80.4%, improving by 1.3 points on average over the previous best method, while maintaining inference speed comparable to recent strong baselines. Project page: https://blaz-r.github.io/changeflow_cd/

Keywords: Change Detection · Rectified Flow · Remote Sensing

1 Introduction

Remote sensing change detection (RSCD) aims to localise changes between two (or more) images of the same geographic region acquired at different times [8, 14]. With the increasing availability of high-resolution remote sensing imagery and advances in deep learning, RSCD has become a key component in applications such as environmental monitoring, land-use mapping, disaster response, and urban development [15, 20, 35, 60]. However, defining exactly what constitutes a change usually requires considering changes at the region level rather than at

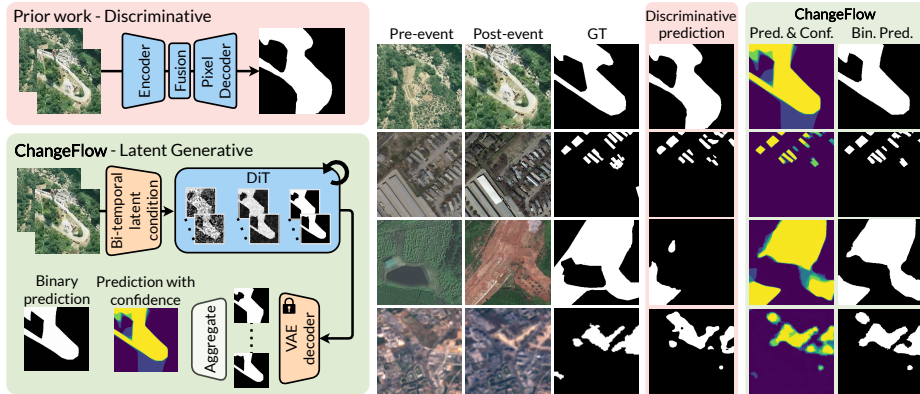


Fig. 1: Unlike discriminative change detection methods, **ChangeFlow** predicts binary change masks through iterative latent generation. This approach enforces global consistency within changed regions and provides better coverage of the changed area. The model inherently enables sampling-based ensembling of predictions, improving results and providing confidence estimation for the change class.

individual pixels, which is inherently ambiguous and based on annotation conventions. Many current change-detection methods cannot effectively capture this, thereby preventing significant advancement in the field.

Most state-of-the-art RSCD approaches follow a *discriminative* paradigm, predicting each pixel independently as *changed* or *unchanged* [3, 8, 13, 44]. While effective, this per-pixel objective provides weak incentives for global mask coherence and becomes a limiting factor since change is defined at the region level. Moreover, standard discriminative methods typically output a single deterministic change mask, which is not well-suited for representing ambiguity and hinders the propagation of confidence to downstream decision-making.

We argue that overcoming this requires a shift from pixel-wise classification to *distribution modelling*. A promising approach here is to use recent generative models, such as rectified flow [33]. They model the distribution of the training data, enabling treating the prediction as a single, coherent concept rather than a set of per-pixel predictions. Additionally, they enable stochastic sampling-based generation of multiple parallel predictions from the same input. Despite this, current generative change detection approaches fail to exploit these concepts, resulting in a significant performance gap compared to discriminative methods [21, 55]. This is largely driven by impractical design choices: current RSCD methods typically operate in pixel space, which is too computationally demanding for iterative generation and unnecessarily difficult for binary masks. Furthermore, they condition the generative process on complex inputs (e.g., auxiliary predictions or elaborate attention mechanisms) that are harder to train, thereby limiting performance.

To address these limitations, we introduce **ChangeFlow**, a generative RSCD framework that reformulates change detection as *change mask synthesis in la-*

latent space using *rectified flow* [33], as illustrated in Figure 1. Specifically, we encode change masks with a pretrained variational autoencoder (VAE) to obtain a compact latent representation. We then train a diffusion transformer (DiT) in rectified flow fashion to transport Gaussian noise to the mask latent space along a straight-line trajectory, enabling efficient sampling with only a few generation steps. We guide (condition) the generative process using features extracted from both input images. Because inference starts from random noise, ChangeFlow naturally supports sampling-based inference without additional training. The samples follow a conditional distribution over change masks given the observed image input, and thereby represent plausible variations of the prediction. Averaging samples reduces prediction variance in a manner similar to classical ensemble methods and naturally yields confidence estimates. The mechanism is particularly effective for change masks, or segmentation maps in general, where the final prediction corresponds to the aggregation of coherent mask hypotheses, an aspect underexplored in segmentation rectified flow models [51] and far less meaningful in current image generation models [33].

In summary, our contributions are threefold: **(i)** we reformulate RSCD as latent-space change mask generation and propose a rectified flow framework that produces globally coherent change masks; **(ii)** we introduce a conditioning strategy based on input feature differences that avoids auxiliary predictors and complex architecture; and **(iii)** we leverage the sampling-based generation inherent to rectified flow models to obtain confidence estimates and effectively fuse predictions, offering a controllable speed–accuracy trade-off by adjusting the number of generation steps and repetitions.

We validate our contributions by evaluating the proposed approach across four standard change detection datasets: SYSU, LEVIR, CLCD, and OSCD, achieving F1 scores of 85.6%, 92.1%, 84.5%, and 59.5%, respectively, substantially outperforming all previous methods on three datasets. This sets a new best average F1 of 80.4% across all four datasets, outperforming the previous-best ChangeDino by 1.3 percentage points.

2 Related work

Remote sensing change detection (RSCD). RSCD has evolved in recent years from pixel-wise differencing and statistical tests to end-to-end deep models [29, 36, 40, 47]. Since early deep models, the field relied on Siamese networks, from convolutional architectures [9, 14, 31], to more recent transformer variants [4, 56, 57], state-space models [10] and diffusion-inspired designs for the backbone [3, 55]. Beyond architectural advances, large-scale pretraining and foundation priors are increasingly important for performance and robustness [13, 30, 44]. Recent work also explores *semantic* change detection, which predicts change together with semantic categories [6, 10, 16, 19]. However, determining whether a change occurred remains the core problem and often generalises beyond fixed label sets. In all settings, the dominant formulation remains discriminative (pixel-wise changed/unchanged classification), which often trades robust change-region

modelling for straightforward supervised training. We instead cast CD as an iterative generative inference problem that explicitly models the distribution of possible change masks, thereby improving mask structure and providing confidence estimates.

Generative models for computer vision tasks. Generative models, particularly diffusion [37] and flow-based [33] formulations, have recently gained traction as powerful tools for visual representation learning. Such models were successfully applied to various fields, such as few-shot counting [49], anomaly detection [18], monocular depth estimation [25], and object detection [12]. Most relevant to our case, it has also been successfully applied to Earth Observation (EO) tasks (e.g., FlowEO [5]) and to general semantic segmentation (e.g., SemFlow [51] and GSS [11]). However, unlike ChangeFlow, such approaches rarely leverage the multiple-samples-based inference that such models offer.

Data synthesis with generative models for change detection. Several works [6, 27, 48, 53, 59] leverage generative models to extend the training set for change detection by synthesising pseudo changes. While effective for increasing data diversity, these approaches treat diffusion solely as an offline generator; change detection is still performed by a separately trained discriminative network. In contrast, we do not rely on synthetic data generation; instead, we formulate CD itself as a generative task.

Diffusion models as feature extractors for change detection. Several methods [3, 22, 23] train diffusion models on remote sensing imagery and use them as feature extractors. The extracted features are then fed to a discriminative head to output a change mask. In contrast, our approach leverages the network’s generative features directly for change-mask prediction, rather than using them solely for feature extraction.

Generative change detection formulations Only a few methods formulate change detection as a generative process. GCD-DDPM [55] conditions diffusion-based generation on the output of another change detection method enhanced with attention. Similarly, SMDNet [21] integrates bi-temporal encodings into a pixel-space DDIM generation process. These methods operate in pixel space, require many generation steps, and rely on complex conditioning mechanisms, which increase computational load and limit performance. In contrast, ChangeFlow utilises a latent rectified flow formulation, avoiding costly pixel-space generation and architecturally complex conditioning schemes, thereby enabling more potent and efficient change-mask generation.

3 Preliminaries

Rectified flow (RF) [33] is a generative framework that maps Gaussian noise $X_0 \sim \mathcal{N}(0, I)$ to a target data distribution $X_1 \sim P_{data}$ via a straight-line trajectory. The intermediate state at any time $t \in [0, 1]$ is defined by linear interpolation:

$$X_t = (1 - t)X_0 + tX_1 . \quad (1)$$

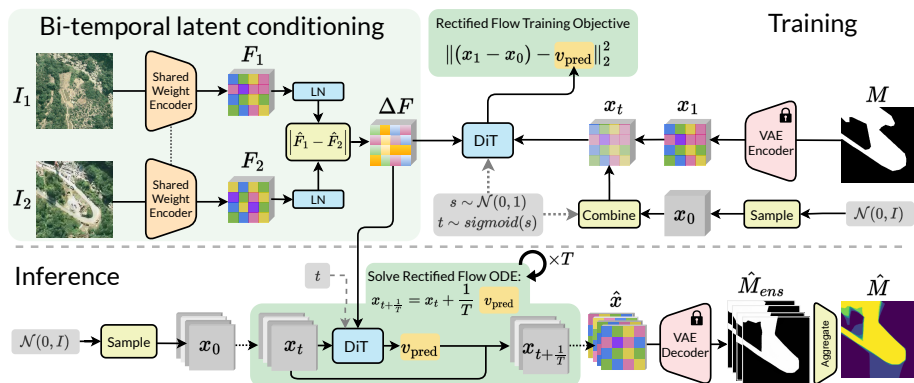


Fig. 2: Up. Training pipeline of ChangeFlow using latent rectified flow conditioned on bi-temporal feature difference. **Down.** During inference, we iteratively generate a change mask by integrating the velocity field. We aggregate multiple samples to form the final prediction and a confidence for the change class from sample agreement.

Because this trajectory has a constant velocity of $(X_1 - X_0)$, a neural network $v_\theta(X_t, t)$ can be trained to predict it by minimising the mean squared error:

$$\min_{\theta} \mathbb{E}_{t, X_0, X_1} [\|(X_1 - X_0) - v_\theta(X_t, t)\|^2] , \quad (2)$$

where t is sampled from $[0, 1]$. During inference, data is generated by integrating the predicted velocity field v_θ starting from an initial noise sample X_0 .

4 ChangeFlow

Recent attempts that use generative modelling for change detection disregard latent formulations, thereby increasing computational complexity. In contrast, we move our modelling process from the pixel to the latent space and use a principled conditioning scheme based on features extracted from a strong pretrained encoder. Given a pair of images, we first extract features using a *Shared Weight Encoder*, and we *condition* the *Diffusion Transformer (DiT)* rectified flow model on the absolute difference of the extracted features. Guided by this conditioning, the model then iteratively generates a latent representation of the corresponding change mask, which is ultimately decoded by the *Variational Autoencoder (VAE)* into a binary change mask. The method is illustrated in Figure 2 and described in detail in the following sections.

4.1 Change detection as latent generative synthesis

Change masks in latent space. To explicitly model the distribution of change masks in latent space and obtain coherent predictions, we formulate change detection as a mask-generation problem. More specifically, we use rectified flow to

generate change masks inside the latent space of a pretrained VAE [26]. While it is known that VAEs efficiently encode RGB images [17, 41], it is unclear whether this holds for binary images (i.e., change masks). To verify this, we perform a simple experiment and report the F1 score and mean absolute error (MAE) in Table 1. We first repeat the binary change mask 3 times along the channel dimension, encode it with the SD-XL [41] VAE, decode the resulting latent, and average the 3 output channels to restore the binary mask. The high F1 score and low MAE indicate that this is indeed feasible and offers potential insights for applications beyond change detection.

Table 1: F1 and mean absolute error (MAE) of binary ground truth masks reconstruction through SD-XL VAE [41].

	SYSU	LEVIR	CLCD	OSCD		SYSU	LEVIR	CLCD	OSCD
F1	99.9	99.3	99.5	99.4	MAE	0.0004	0.0007	0.0006	0.0006

Change mask rectified flow. Let $M \in \{0, 1\}^{H \times W}$ denote the binary ground-truth change mask where H and W are the mask dimensions and \mathcal{V} is a pretrained VAE encoder $\mathcal{V}: \mathbb{R}^3 \rightarrow \mathbb{R}^d$ (in our case SD-XL [41] VAE). As described in the previous section, we can then encode the change mask with \mathcal{V} as:

$$x_1 = \mathcal{V}(\{M, M, M\}), \quad x_1 \in \mathbb{R}^{h \times w \times d}, \quad (3)$$

where $\{\cdot, \cdot, \cdot\}$ indicates value repeating in channel dimension. This yields a compact latent representation $h < H, w < W$.

During training, we sample Gaussian noise in the same shape as the latent space to obtain an *initial state* x_0 :

$$x_0 \sim \mathcal{N}(0, I), \quad x_0 \in \mathbb{R}^{h \times w \times d}, \quad (4)$$

which we use to construct an interpolated latent (i.e., an intermediate step along the straight trajectory) representation at a specified time step t :

$$x_t = (1 - t)x_0 + tx_1. \quad (5)$$

Previous work [17] has shown the importance of selecting the correct distribution for sampling timesteps during training. Therefore, we sample timesteps in a logit-normal fashion, which emphasises learning at the critical point where $t = 0.5$:

$$t \in [0, 1]; \quad t = \text{sigmoid}(s); \quad s \sim \mathcal{N}(0, 1). \quad (6)$$

This represents the most ambiguous point in time at which the levels of noise and signal are balanced, with trajectories overlapping the most, and the model must learn to rectify the field (refer to [33] for more details).

To guide the network from initial noise to the final mask latent space, we prepare a bi-temporal latent conditioning signal ΔF , which we will explain at the end of this subsection. We concatenate it with x_t in the channel dimension

and feed the resulting vector to the model. The rectified flow vector field is then parametrised using a DiT [39]-based network \mathcal{M}_θ :

$$v_{\text{pred}} = \mathcal{M}_\theta([x_t, \Delta F], t) . \quad (7)$$

We train the network using the standard MSE loss for rectified flows [33]:

$$\mathcal{L}_{\text{RF}} = \|(x_1 - x_0) - v_{\text{pred}}\|_2^2 . \quad (8)$$

This means that there is no explicit per-pixel objective; the model learns the velocity field at a specific time step (i.e., at a specific location along the straight trajectory). The process is also depicted in the top of Figure 2.

Change mask generation guidance. To create a conditioning signal used to guide the generation process, we first extract high-level latent features from an image pair (I_1, I_2) using a pretrained encoder Φ with shared weights:

$$F_1 = \Phi(I_1), \quad F_2 = \Phi(I_2), \quad F_1, F_2 \in \mathbb{R}^{h' \times w' \times c} . \quad (9)$$

To remain agnostic to temporal ordering and feature magnitude, we construct the conditioning signal as the absolute difference of the layer normalised (LayerNorm [1] - LN) feature maps:

$$\Delta F = |\text{LN}(F_1) - \text{LN}(F_2)| . \quad (10)$$

The process is also illustrated in the top-left of Figure 2. Unlike previous generative change detection works [21, 55], this approach avoids complex auxiliary methods and attention-based conditioning, offering an efficient latent design that enables the model to learn optimal latent conditioning for the task.

4.2 Inference via rectified flow integration

At inference time, given a pair of images I_1 and I_2 , we compute ΔF (explained in the previous section) and sample an initial noise:

$$x_0 \sim \mathcal{N}(0, I) . \quad (11)$$

The change mask latent is then generated by solving the rectified flow ordinary differential equation (ODE) using Euler integration over equally spaced T steps:

$$x_{t+\frac{1}{T}} = x_t + \frac{1}{T} \mathcal{M}_\theta([x_t, \Delta F], t) . \quad (12)$$

The final latent $\hat{x} = x_T$ is decoded into a binary RGB change mask using the pretrained VAE decoder \mathcal{V}^{-1} :

$$\hat{M}_{\text{RGB}} = \mathcal{V}^{-1}(\hat{x}) . \quad (13)$$

To obtain the final single-channel binary mask, the prediction is averaged across the RGB channels, yielding \hat{M} . The entire inference process is depicted at the bottom of Figure 2. By using the rectified flow formulation, we allow for a flexible number of time steps at inference, which can be *freely adjusted after training* based on available computing resources.

4.3 Ensembling and confidence

Our formulation enables sampling-based inference without additional training, thereby facilitating the ensembling of multiple predictions and improving performance. The rectified flow model implicitly defines a conditional distribution [33] over change masks by marginalising latent noise, i.e., $p(M | \Delta F) = \int p(M | \Delta F, x_0) p(x_0) dx$. In practice, this marginalization is approximated via Monte Carlo sampling by generating ensemble masks $\hat{M}_{ens}^{(i)}, i \in \{1, \dots, N\}$ starting from different initial noise x_0^i and aggregating them into a joint prediction \hat{M} (e.g., via a mean or majority vote). Since masks are binary in our case, we use simple averaging aggregation.

This process also provides a clear *confidence signal* regarding the change class. The per-pixel mask mean reflects agreement across hypotheses, with lower values in ambiguous changed regions and higher values where predictions consistently coincide. In contrast, obtaining such confidence from standard discriminative models typically requires additional mechanisms (e.g., confidence heads or losses [50, 54]), rather than arising as an inherent property of the model.

5 Results

Implementation details. We use DINOv3 [46] ViT-L as the encoder and extract features from its final layer. For mask encoding, we adopt the VAE from SD-XL [41]. To spatially align the encoder and VAE latents, we apply bicubic interpolation to the conditioning tensor. Each inference involves 10 steps (i.e., $T = 10$). We generate an ensemble of 5 predictions and fit the standard CD metrics by binarising: a pixel is marked changed if at least 2 predictions mark it as such. Input images are cropped to 256×256 pixels and augmented with random flips and rotations during training. We train using the Muon [24] optimiser, with an initial learning rate of 10^{-4} for DiT and $5 \cdot 10^{-5}$ for the encoder, and a cosine scheduler without restarts. Training lasts 300 epochs with a batch size of 32 on an NVIDIA A100 GPU. Additional details are in the Supplementary.

Evaluation metrics and datasets. We evaluate change detection performance using *binary* precision, recall, and F1, considering only *change class* [4, 8, 14, 44] on the model from the final epoch. For robust evaluation, we benchmark on four change detection datasets covering *diverse locations, sensors, and ground sampling distances*, and spanning *diverse change types*, including *building, urban, and cropland* changes, as well as changes resulting from *natural disasters*. *SYSU* [45] covers various change types, from buildings and vegetation to sea changes. *LEVIR* [9] focuses on building changes, while *CLCD* [32] captures only changes that happen on croplands. *OSCD* [15] is a low-resolution global Sentinel-2 dataset covering urban changes. Models are trained on a dedicated training set and evaluated on the test set. Additional details are in the Supplementary.

5.1 Main results

Change detection methods. We evaluate ChangeFlow against a range of change detection methods, including discriminative architectures FCSDiff [14], ChangeFormer [4], SwinSUNet [57], GFM [34], BiFA [58], MaskCD [56], ChangeMamba [10], MTP [52], HySCDG [6], BTC [44] and ChangeDINO [13]. We also compare to the generative GCD-DDPM [55], as well as diffusion-based discriminative methods DDPM-CD [3] and SatDiFuser [22]. Implementation details are in the Supplementary. ChangeDINO [13] in particular represents the current state-of-the-art and uses the same DINOv3 [46] backbone as our proposed method, ChangeFlow. We summarize quantitative results across all datasets and methods in Table 2. Extended results are in the Supplementary.

Table 2: Comparison of our proposed method, ChangeFlow, to the state-of-the-art on four different datasets. We mark **first**, **second**, and **third** place results. FPS was benchmarked on an NVIDIA A100 using the protocol described in the Supplementary.

	FPS	Param.	SYSU	LEVIR	CLCD	OSCD	<i>Avg</i>
	[img/s]	[M]	F1	F1	F1	F1	F1
FC-Siam-Diff [14] _{ICIP18}	170.1±0.0	1.4	70.8	81.8	54.1	39.4	61.5
ChFormer [4] _{IGARSS22}	36.2±0.0	41.0	77.9	89.5	60.8	48.1	69.1
SwinSUNet [57] _{TGRS22}	33.1±0.1	43.6	76.6	89.3	75.8	52.8	73.6
GFM [34] _{CVPR23}	44.9±0.1	120.5	81.2	89.8	77.5	54.1	75.7
GCD-DDPM [55] _{TGRS24}	0.02±0.0	131.9	64.5	80.7	46.9	7.0	49.8
BiFA [58] _{TGRS24}	32.2±0.0	9.9	83.8	89.5	74.5	37.4	71.3
MaskCD [56] _{TGRS24}	6.5±0.0	107.4	83.8	90.3	76.6	34.7	71.4
ChMamba [10] _{TGRS24}	14.4±0.0	92.4	81.5	91.8	80.3	45.8	74.9
MTP [52] _{JSTARS24}	31.2±0.0	107.8	81.3	91.7	80.3	52.8	76.5
HySCDG [6] _{CVPR25}	41.0±0.1	65.1	78.7	91.1	64.3	53.6	71.9
DDPM-CD [3] _{WACV25}	4.6±0.0	437.5	80.5	90.9	71.4	37.1	70.0
SatDiFuser [22] _{ICCV25}	1.8±0.0	1413.6	82.0	90.2	79.1	55.2	76.6
BTC [44] _{TGRS25}	32.4±0.0	120.1	82.4	91.5	80.9	54.3	77.3
ChangeDINO [13] _{arXiv25}	8.9±0.5	311.1	83.9	92.2	81.4	58.8	79.1
ChangeFlow _(10step, 5rep)	8.1±0.0	403.3	85.6	92.1	84.5	59.5	80.4
ChangeFlow _(10step, 1rep)	15.3 ±0.0	403.3	83.9	92.0	84.4	57.5	79.5
ChangeFlow _(1step, 5rep)	18.2 ±0.0	403.3	85.6	92.0	84.5	59.0	80.3

Quantitative results. ChangeFlow achieves the best average F1 score of 80.4%, which is 1.3 points higher than the previous best method, and establishes a new state-of-the-art result on SYSU, CLCD, and OSCD with F1 scores of 85.6%, 84.5% and 59.5%, respectively. These datasets all contain challenging, highly semantic changed regions where mask coherence is important. On LEVIR, our method remains competitive and is within 0.1 percentage points of the best competing method. Despite iterative inference, ChangeFlow, with 10 steps and 5 ensemble predictions, achieves a similar throughput to ChangeDINO. Moreover, a single-step variant of our method with 5 ensemble predictions (1step, 5rep)

maintains nearly the same accuracy while substantially increasing throughput, demonstrating a good trade-off between speed and accuracy.

Comparison to diffusion-based methods. ChangeFlow outperforms all prior discriminative approaches that use diffusion models as feature extractors, such as DDPM-CD, by 10.3 p.p., and the recent SatDiFuser foundation model by 4.8 p.p. It also substantially exceeds the pixel-space generative baseline, GCD-DDPM, by 30.5 p.p., while being almost 3 orders of magnitude faster in inference. This shows that we have substantially improved both performance and speed compared to the previous generative attempt.

Qualitative results. In Figure 3, we show a qualitative comparison between evaluated methods. Compared with ChangeDINO, our method reduces missed detections in homogeneous regions, consistent with its coherent mask-generation behaviour. Compared to DDPM-CD, which uses diffusion primarily as a feature extractor, ChangeFlow better recovers complete change regions and reduces both false positives and false negatives. We also include MaskCD [56], which extends pixel-wise classification by predicting mask instances. Still, MaskCD’s predictions remain fragmented across diverse change types, whereas ChangeFlow directly generates a globally consistent mask. More qualitative results, including failure cases, are in the Supplementary.




Fig. 3: Qualitative comparison of competing methods. The pair of considered images is shown in the first and second columns, followed by the ground truth mask and predictions for the related methods and our method. False positives are marked in red and false negatives in blue.

Generation visualization. To better illustrate our generative process, we visualise intermediate generation steps over $T=10$ inference steps in Figure 4 (all but last column). The model rapidly forms a coarse change region and then refines boundaries in later steps. Compared to the pixel-space diffusion CD method GCD-DPPM [55] that requires 1000 steps, ChangeFlow achieves strong results with a much smaller number of steps, enabling efficient inference in practice.

Coherence and confidence analysis. To better estimate model confidence for the change class, we use sample agreement across repetitions: pixels that are consistently predicted as changed across hypotheses represent higher confidence. Visually, this appears as lower intensity in ambiguous regions (Figure 4 last column). By performing thresholding based on these predictions, we can practically adjust the operating point (e.g., favouring precision or recall) without modifying the model. Extended analysis with more images is in the Supplementary.




Fig. 4: Visualisation of intermediate steps (stride 2 for visualisation purposes) in the latent generative mask prediction (all but last column) and confidence obtained from an ensemble of predictions (last column). ChangeFlow performs iterative prediction from pure noise to a binary mask, as explained in Section 4.1. Here, we decode the intermediate latent representation into a binary mask at each step. The final ensemble of predictions is illustrated by high-confidence regions being brighter in colour.

One of the great advantages of ChangeFlow’s formulation is the inherent global prediction coherence. To quantitatively evaluate this, we assess structural consistency by calculating the error from the expected ground-truth number of connected components and hole counts (reported as Δ) across 4 datasets. Implementation details are in the Supplementary. Figure 5 shows that ChangeFlow yields low structural error, indicating the fewest spurious holes and few incorrectly fragmented components.

To better understand the generation process, we calculate these metrics at intermediate generation steps and report the results in Figure 6. The model produces the coherent region from the initial noise relatively quickly (low hole and connected component (CC) errors from step 5 onwards). Initially, due to noisy latent, the prediction is just a single random component with many holes. In the next 3 steps, it breaks into many small components, and by step 5, it already forms a coherent region (see Figure 4 for visualisation).




Fig. 5: Coherence measured as hole count error ($\Delta \#Holes$) and connected component count error ($\Delta \text{ Connected Components}$) averaged over 4 datasets. Lower is better.




Fig. 6: Coherence measured in connected component count error (ΔCC) and hole count error ($\Delta Holes$) averaged over 4 datasets for 10 generation steps. Lower is better.

5.2 Ablation study

In this section, we isolate the impact of our contributions by ablating key design choices. Implementation details and more ablations are in the Supplementary.

Table 3: Ablation of design choices when constructing conditioning (Section 4.1).

ID		SYSU	LEVIR	CLCD	OSCD	Avg
		F1	F1	F1	F1	F1
<i>Ours</i>	Abs. diff. with LayerNorm	85.6	92.1	84.5	59.5	80.4
<i>Sub</i>	Abs. diff. \rightarrow Difference	81.2	91.7	82.2	35.6	72.7
<i>Concat</i>	Abs. diff. \rightarrow Concat	77.8	91.4	80.9	21.9	68.0
<i>L2Norm</i>	LayerNorm \rightarrow L2 norm	85.2	92.0	84.2	57.2	79.6
<i>NoNorm</i>	No normalisation	81.8	91.4	77.6	56.4	76.8

Conditioning ablations. In Table 3, we ablate various design choices for conditioning vectors that guide the generation process. Using the absolute feature difference with LayerNorm [1] normalisation consistently delivers the best overall performance. Replacing absolute difference with signed subtraction (*Sub*) reduces average F1 by 7.7 p.p., as the conditioning becomes sensitive to temporal order. The performance degrades further when features are concatenated rather than fused elementwise (*Concat*), resulting in a 12.4 p.p. drop, particularly on OSCD, suggesting that non-structured conditioning is harder to utilise effectively. Using LayerNorm consistently benefits performance: removing it reduces performance by 3.6 p.p. (*NoNorm*), while the choice between LayerNorm and simple channel-wise L2 normalisation has a smaller effect (*L2Norm*). This suggests that the normalisation type is less critical than its presence. These ablations also underscore our motivation: keeping the conditioning simple but structured yields the best performance, without the need for unnecessarily complex design.

Encoder ablations. In Table 4, we show that DINOv3 on average provides the strongest features for ChangeFlow. The satellite-pretrained variant (DINOv3 Sat.) and DINOv2 perform worse, most likely due to reduced generalisation from smaller pretraining corpora. RADIO yields solid results, with version 4 performing considerably better than 2.5, but does not surpass plain DINOv3.

Table 4: Ablation of shared weight encoder selection.

	SYSU F1	LEVIR F1	CLCD F1	OSCD F1	<i>Avg</i> F1
DINOv3	85.6	92.1	84.5	59.5	80.4
DINOv3 Sat. [46]	83.6	91.6	80.7	59.4	78.9
DINOv2 [38]	78.4	91.6	78.4	54.8	75.8
RADIO 2.5 [42]	80.7	91.3	78.7	58.8	77.4
RADIO 4 [43]	84.2	91.9	82.8	57.8	79.2

Table 5: Ablation of VAEs used for encoding labels during training and decoding of binary masks during inference.

	Frozen Decoder	SYSU F1	LEVIR F1	CLCD F1	OSCD F1	<i>Avg</i> F1
SD-XL VAE [41]	✓	85.6	92.1	84.5	59.5	80.4
SD 3.5 [17]	✓	84.4	91.6	84.4	56.7	79.3
Z-Image [7]	✓	85.2	91.7	83.2	57.8	79.5
Flux.1-dev [28]	✓	84.7	91.7	82.1	57.2	78.9
SD-XL VAE-Finetuned		84.4	92.1	83.6	57.2	79.3
SD-XL VAE-Finetuned (Pixel loss only)		84.1	92.0	81.9	36.3	73.6
CNN Decoder		84.1	89.4	83.4	54.7	77.9
CNN Decoder (Pixel loss only)		82.7	92.4	82.0	43.9	75.2

VAE ablations. In Table 5, we compares VAEs for mask encoding and decoding. Among pretrained VAEs, the SD-XL VAE [41], with a latent dimension of 4, achieves the best average performance. In contrast, the VAEs with latent dimension of 16 (SD 3.5 [17], Z-image [7], and Flux. 1-dev [28]) are consistently slightly worse. A plausible explanation is that the higher latent dimensionality increases the difficulty of learning a well-conditioned rectified flow transport for sparse binary masks. Importantly, all pretrained VAEs remain competitive overall, suggesting that off-the-shelf VAE latents provide practical and effective representations for change-mask generation.

We also investigate whether adapting the decoder improves results. Jointly finetuning the SD-XL VAE decoder with an auxiliary Dice loss on mask does not yield gains, and replacing the rectified flow objective with a pure pixel-level loss

(gradients are also passed through DiT) degrades performance, indicating that the generative training signal is a key contributor to the final accuracy, similar to observation from [2].

Finally, we replace the VAE decoder with a lightweight UNet-like CNN decoder. This alternative is generally weaker on average, both when trained with the rectified flow objective and when trained with a pixel-only loss. Nevertheless, on LEVIR, the pixel-only CNN variant performs best, which we attribute to the sharper delineation of small building boundaries in this dataset.

Inference steps and repetition ablation. Our default model formulation uses $T = 10$ steps and 5 predictions (repetitions) in the ensemble. Figure 7 shows that increasing repetitions at fixed steps yields more consistent gains than increasing T beyond a small number of steps at fixed repetitions, while both increase runtime. This provides a controllable speed–accuracy trade-off at inference time.




Fig. 7: Impact of number of sampling steps (at fixed rate of 5 repetitions) and inference repetitions (at fixed 10 sampling steps). Change detection performance is reported on the left y-axis, measured as average F1 across 4 datasets, while inference speed is reported on the right y-axis as frames per second (FPS; protocol in the Supplementary).

Limitations and future work. ChangeFlow currently relies on a generic pre-trained VAE that is not explicitly designed for the binary, sparse structure of change masks. Even though we show that the data loss is minimal in Table 1, designing VAEs tailored to binary latent spaces may further improve boundary precision and generation stability. For real-world deployment, further improving inference speed is another possibility, which aligns with general trends in generative modelling. Finally, the generative formulation naturally opens the possibility of incorporating textual guidance into the latent inference process, enabling more flexible and potentially open-vocabulary change detection.

6 Conclusion

We have introduced ChangeFlow, a latent generative method for remote sensing change detection (RSCD) based on rectified flow. Unlike prior generative RSCD methods that operate in pixel space and use complex conditioning schemes, ChangeFlow generates change masks in the latent space and uses a simpler conditioning signal derived from bi-temporal feature differences. Our generative

formulation enables efficient inference with a small number of generation steps and naturally supports sampling-based prediction: aggregating multiple samples improves robustness, and sample agreement provides a practical estimation for the change class’s confidence. Evaluation on four benchmarks (SYSU, LEVIR, CLCD, OSCD) yields strong results, with an average F1 score of 80.4%, 1.3 points higher than the previous best-performing method. Gains are most pronounced on data with low-resolution imagery (e.g., Sentinel 2) and highly semantic region-level changes. Our results indicate that latent-space generative inference is a conceptual step forward from purely discriminative methods for producing coherent change masks in challenging RSCD settings, and motivate further exploration of flow-based mask synthesis for other dense prediction tasks.

Acknowledgements

This work was in part supported by the ARIS research projects GC-0006 (GeoAI) and J2-60045 (RoDEO), research programme P2-0214, and the supercomputing network SLING (ARNES, EuroHPC Vega).

References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer Normalization. arXiv preprint arXiv:1607.06450 (2016) [7](#), [12](#), [4](#), [13](#)
2. Bagchi, A., Bao, Z., Wang, Y.X., Tokmakov, P., Hebert, M.: ReferEverything: Towards Segmenting Everything We Can Speak of in Videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23221–23231 (2025) [14](#)
3. Bandara, W.G.C., Nair, N.G., Patel, V.: DDPM-CD: Denoising Diffusion Probabilistic Models as Feature Extractors for Remote Sensing Change Detection. In: IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5250–5262 (2025) [2](#), [3](#), [4](#), [9](#), [11](#)
4. Bandara, W.G.C., Patel, V.M.: A Transformer-Based Siamese Network for Change Detection. In: IEEE International Geoscience and Remote Sensing Symposium. pp. 207–210. IEEE (2022) [3](#), [8](#), [9](#), [2](#), [11](#), [16](#)
5. Bellier, G.L., Audebert, N.: FlowEO: Generative Unsupervised Domain Adaptation for Earth Observation. arXiv preprint arXiv:2512.05140 (2025) [4](#)
6. Benidir, Y., Gonthier, N., Mallet, C.: The Change You Want To Detect: Semantic Change Detection In Earth Observation With Hybrid Data Generation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 2204–2214 (2025) [3](#), [4](#), [9](#), [2](#), [11](#), [16](#)
7. Cai, H., Cao, S., Du, R., Gao, P., Hoi, S., Huang, S., Hou, Z., Jiang, D., Jin, X., Li, L., et al.: Z-Image: An Efficient Image Generation Foundation Model With Single-Stream Diffusion Transformer. arXiv preprint arXiv:2511.22699 (2025) [13](#)
8. Chen, H., Qi, Z., Shi, Z.: Remote Sensing Image Change Detection With Transformers. IEEE Transactions on Geoscience and Remote Sensing **60**, 1–14 (2021) [1](#), [2](#), [8](#)
9. Chen, H., Shi, Z.: A Spatial-Temporal Attention-Based Method and A New Dataset for Remote Sensing Image Change Detection. Remote Sensing **12**, 1662 (2020) [3](#), [8](#), [1](#)

10. Chen, H., Song, J., Han, C., Xia, J., Yokoya, N.: ChangeMamba: Remote Sensing Change Detection With Spatiotemporal State Space Model. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–20 (2024) [3](#), [9](#), [2](#), [11](#), [16](#)
11. Chen, J., Lu, J., Zhu, X., Zhang, L.: Generative Semantic Segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7111–7120 (2023) [4](#)
12. Chen, S., Sun, P., Song, Y., Luo, P.: DiffusionDet: Diffusion Model for Object Detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 19830–19843 (2023) [4](#)
13. Cheng, C.H., Hsu, C.C.: ChangeDino: DINOv3-Driven Building Change Detection in Optical Remote Sensing Imagery. *arXiv preprint arXiv:2511.16322* (2025) [2](#), [3](#), [9](#), [11](#), [17](#)
14. Daudt, R.C., Le Saux, B., Boulch, A.: Fully Convolutional Siamese Networks for Change Detection. In: *IEEE International Conference on Image Processing*. pp. 4063–4067. *IEEE* (2018) [1](#), [3](#), [8](#), [9](#), [2](#), [11](#), [16](#)
15. Daudt, R.C., Le Saux, B., Boulch, A., Gousseau, Y.: Urban Change Detection for Multispectral Earth Observation Using Convolutional Neural Networks. In: *IEEE International Geoscience and Remote Sensing Symposium*. pp. 2115–2118. *IEEE* (2018) [1](#), [8](#)
16. Ding, L., Zhang, J., Guo, H., Zhang, K., Liu, B., Bruzzone, L.: Joint Spatio-Temporal Modeling for Semantic Change Detection in Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–14 (2024) [3](#)
17. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In: *Forty-first International Conference on Machine Learning* (2024) [6](#), [13](#), [3](#)
18. Fučka, M., Zavrtnik, V., Skočaj, D.: Transfusion—a Transparency-Based Diffusion Model for Anomaly Detection. In: *European Conference on Computer Vision*. pp. 91–108. *Springer* (2024) [4](#)
19. Guo, H., Liu, C., Zhang, H., Chen, B., Zou, Z., Shi, Z.: TaCo: Capturing Spatio-Temporal Semantic Consistency in Remote Sensing Change Detection. *arXiv preprint arXiv:2511.20306* (2025) [3](#)
20. Hänsch, R., Chaurasia, M.A.: Earth Observation and Machine Learning for Climate Change. In: *IEEE International Geoscience and Remote Sensing Symposium*. pp. 1676–1682. *IEEE* (2024) [1](#)
21. Jia, J., Lee, G., Wang, Z., Lyu, Z., He, Y.: Siamese Meets Diffusion Network: Smdnet for Enhanced Change Detection in High-Resolution Rs Imagery. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **17**, 8189–8202 (2024) [2](#), [4](#), [7](#)
22. Jia, Y., Marsocci, V., Gong, Z., Yang, X., Vergauwen, M., Nascetti, A.: Can Generative Geospatial Diffusion Models Excel as Discriminative Geospatial Foundation Models? In: *International Conference on Computer Vision (ICCV)* (2025) [4](#), [9](#), [2](#), [11](#), [16](#)
23. Jiang, F., Huo, X., Zhang, M., Gong, M., Pu, Y., Zhou, Y., Zhao, W., Guan, Z.: D3PM: Dual-Stream Denoising Diffusion Probabilistic Model for Change Detection in Multimodal Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing* (2025) [4](#)
24. Jordan, K., Jin, Y., Boza, V., Jiacheng, Y., Cesista, F., Newhouse, L., Bernstein, J.: Muon: An Optimizer for Hidden Layers in Neural Networks (2024), <https://kellerjordan.github.io/posts/muon/> [8](#)

25. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daudt, R.C., Schindler, K.: Repurposing Diffusion-Based Image Generators for Monocular Depth Estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9492–9502 (2024) [4](#)
26. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings (2014) [6](#)
27. Korkmaz, Y., Paranjape, J.N., de Melo, C.M., Patel, V.M.: Referring Change Detection in Remote Sensing Imagery. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2026) [4](#)
28. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux> (2024) [13](#)
29. Le Saux, B., Randrianarivo, H.: Urban Change Detection in SAR Images by Interactive Learning. In: IEEE International Geoscience and Remote Sensing Symposium. pp. 3990–3993. IEEE (2013) [3](#)
30. Li, K., Cao, X., Meng, D.: A New Learning Paradigm for Foundation Model-Based Remote-Sensing Change Detection. IEEE Transactions on Geoscience and Remote Sensing **62**, 1–12 (2024) [3](#)
31. Li, Z., Tang, C., Liu, X., Zhang, W., Dou, J., Wang, L., Zomaya, A.Y.: Lightweight Remote Sensing Change Detection With Progressive Feature Aggregation and Supervised Attention. IEEE Transactions on Geoscience and Remote Sensing **61**, 1–12 (2023) [3](#)
32. Liu, M., Chai, Z., Deng, H., Liu, R.: A CNN-Transformer Network With Multiscale Context Aggregation for Fine-Grained Cropland Change Detection. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing **15**, 4297–4306 (2022) [8](#), [1](#)
33. Liu, X., Gong, C., Liu, Q.: Flow Straight and Fast: Learning to Generate and Transfer Data With Rectified Flow. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1–5, 2023 (2023) [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
34. Mendieta, M., Han, B., Shi, X., Zhu, Y., Chen, C.: Towards Geospatial Foundation Models via Continual Pretraining. In: IEEE/CVF International Conference on Computer Vision. pp. 16806–16816 (2023) [9](#), [2](#), [11](#), [16](#)
35. Meneses III, S., Blanco, A.: Rapid Mapping and Assessment of Damages Due to Typhoon Rai Using Sentinel-1 Synthetic Aperture Radar Data. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences **43**, 1139–1146 (2022) [1](#)
36. Metzger, N., Türkoglu, M.Ö., Daudt, R.C., Wegner, J.D., Schindler, K.: Urban Change Forecasting From Satellite Images. PFG–Journal of Photogrammetry, Remote Sensing and Geoinformation Science **91**(6), 443–452 (2023) [3](#)
37. Nichol, A.Q., Dhariwal, P.: Improved Denoising Diffusion Probabilistic Models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021) [4](#), [17](#)
38. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., HAZIZA, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning Robust Visual Features Without Supervision. Transactions on Machine Learning Research (2024) [13](#)
39. Peebles, W.S., Xie, S.: Scalable Diffusion Models With Transformers. 2023 Ieee. In: CVF International Conference on Computer Vision (ICCV). vol. 4172 (2022) [7](#)

40. Peng, D., Liu, X., Zhang, Y., Guan, H., Li, Y., Bruzzone, L.: Deep Learning Change Detection Techniques for Optical Remote Sensing Imagery: Status, Perspectives and Challenges. *International Journal of Applied Earth Observation and Geoinformation* **136**, 104282 (2025) [3](#)
41. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In: *The Twelfth International Conference on Learning Representations* (2024) [6](#), [8](#), [13](#), [12](#)
42. Ranzinger, M., Heinrich, G., Kautz, J., Molchanov, P.: Am-RADIO: Agglomerative Vision Foundation Model Reduce All Domains Into One. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 12490–12500 (June 2024) [13](#)
43. Ranzinger, M., Heinrich, G., McCarthy, C., Kautz, J., Tao, A., Catanzaro, B., Molchanov, P.: C-RADIOv4 (tech Report). arXiv preprint arXiv:2601.17237 (2026) [13](#)
44. Rolih, B., Fučka, M., Wolf, F., Čehovin Zajc, L.: Be the Change You Want to See: Revisiting Remote Sensing Change Detection Practices. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–11 (2025) [2](#), [3](#), [8](#), [9](#), [11](#), [12](#), [14](#), [15](#), [16](#), [17](#)
45. Shi, Q., Liu, M., Li, S., Liu, X., Wang, F., Zhang, L.: A Deeply Supervised Attention Metric-Based Network and an Open Aerial Image Dataset for Remote Sensing Change Detection. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–16 (2022) [8](#), [1](#)
46. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: DINOv3. arXiv preprint arXiv:2508.10104 (2025) [8](#), [9](#), [13](#)
47. Singh, A.: Review Article Digital Change Detection Techniques Using Remotely-Sensed Data. *International journal of remote sensing* **10**(6), 989–1003 (1989) [3](#)
48. Song, J., Chen, H., Yokoya, N.: Syntheworld: A Large-Scale Synthetic Dataset for Land Cover Mapping and Building Change Detection. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 8287–8296 (2024) [4](#)
49. Šuštar, G., Pelhan, J., Lukežič, A., Kristan, M.: CoDi—an Exemplar-Conditioned Diffusion Model for Low-Shot Counting. arXiv preprint arXiv:2512.20153 (2025) [4](#)
50. Wan, S., Wu, T.Y., Wong, W.H., Lee, C.Y.: ConfNet: Predict With Confidence. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 2921–2925. IEEE (2018) [8](#)
51. Wang, C., Li, X., Qi, L., Ding, H., Tong, Y., Yang, M.H.: SemFlow: Binding Semantic Segmentation and Image Synthesis via Rectified Flow. *Advances in Neural Information Processing Systems* **37**, 138981–139001 (2024) [3](#), [4](#)
52. Wang, D., Zhang, J., Xu, M., Liu, L., Wang, D., Gao, E., Han, C., Guo, H., Du, B., Tao, D., et al.: MTP: Advancing Remote Sensing Foundation Model via Multi-Task Pretraining. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2024) [9](#), [2](#), [11](#), [14](#), [16](#)
53. Wang, J.X., Li, T., Chen, S.B., Gu, C.J., You, Z.H., Luo, B.: Diffusion Models and Pseudo-Change: A Transfer Learning-Based Change Detection in Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing* (2024) [4](#)
54. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: DUS3R: Geometric 3D Vision Made Easy. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20697–20709 (2024) [8](#)

55. Wen, Y., Ma, X., Zhang, X., Pun, M.O.: Gcd-ddpm: A Generative Change Detection Model Based on Difference-Feature-Guided DDPM. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–16 (2024) [2](#), [3](#), [4](#), [7](#), [9](#), [10](#), [11](#), [16](#)
56. Yu, W., Zhang, X., Das, S., Zhu, X.X., Ghamisi, P.: MaskCD: A Remote Sensing Change Detection Network Based on Mask Classification. *IEEE Transactions on Geoscience and Remote Sensing* (2024) [3](#), [9](#), [10](#), [2](#), [11](#), [16](#)
57. Zhang, C., Wang, L., Cheng, S., Li, Y.: SwinSUNet: Pure Transformer Network for Remote Sensing Image Change Detection. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–13 (2022) [3](#), [9](#), [2](#), [11](#), [16](#)
58. Zhang, H., Chen, H., Zhou, C., Chen, K., Liu, C., Zou, Z., Shi, Z.: Bifa: Remote Sensing Image Change Detection With Bitemporal Feature Alignment. *IEEE Transactions on Geoscience and Remote Sensing* (2024) [9](#), [2](#), [11](#), [16](#)
59. Zheng, Z., Ermon, S., Kim, D., Zhang, L., Zhong, Y.: Changen2: Multi-Temporal Remote Sensing Generative Change Foundation Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **47**(2), 725–741 (2025) [4](#)
60. Zhu, Z., Qiu, S., Ye, S.: Remote Sensing of Land Change: A Multifaceted Perspective. *Remote Sensing of Environment* **282**, 113266 (2022) [1](#)

ChangeFlow - Latent Rectified Flow for Change Detection in Remote Sensing

Supplementary Material

In this Appendix, we provide additional details that extend beyond the scope of the main manuscript. The Appendix is organised as follows:

- **Extended dataset details** in Section [A](#).
- **Extended results** with additional metrics (precision and recall), additional ablations, as well as an additional confidence evaluation in Section [B](#).
- **Additional qualitative results**, including failure cases, VAE mask reconstruction, intermediate step generation and confidence visualisation in Section [C](#).
- **Computational efficiency** protocol, extended results, and discussion in Section [D](#).
- **Extended implementation details** for our model and training, our ablations and analyses, and related methods in Section [E](#).

A Extended dataset details

	Acquisition	Resolution	Change Type	Interval	Region	Image count train\val\test	Patch	Changed Pixels	Unchanged Pixels
SYSU [45]	Aerial	0.5m	Building, urban, groundwork, road, vegetation, sea	2007-2014	Hong Kong	12000 4000 4000	256 × 256	21.8 %	78.2 %
LEVIR [9]	Google Earth satellite	0.5m	Building	2002-2018	20 regions in US	7120 1024 2048	256 × 256	4.7 %	95.3 %
CLCD [32]	Satellite (Gaofen-2)	0.5m-2m	Multiple types limited to croplands	2017-2019	Guangdong, China	1440 480 480	256 × 256	7.6 %	92.4 %
OSCD [15]	Satellite (Sentinel-2)	10m	Urban	2015-2018	24 regions worldwide	827 - 385	96 × 96	3.2 %	96.8 %

Table 1: Additional details for the datasets used in the paper.

Additional dataset details are provided in [Table 1](#). Our benchmarks span diverse change scenarios, including building and urban expansion, as well as changes limited to croplands. They also vary substantially in ground sampling distance (GSD), acquisition sensor, and scale, ranging from a few hundred to several thousand image pairs. This diversity strengthens the robustness and generality of our conclusions.

A recurring challenge in RSCD is severe class imbalance: changed pixels typically constitute less than 10% of all pixels. SYSU is an exception, exhibiting a higher change ratio due to larger changed regions.

A.1 Data implementation details

Dataset splits. Official train and test splits are used for OSCD, SYSU, CLCD, and LEVIR to ensure full reproducibility and fair comparison. We also use a validation set (for optimal threshold computation) from SYSU, CLCD, and LEVIR.

The HuggingFace public source for the data used is as follows:

- SYSU: [ericyu: SYSU_CD](#)
- LEVIR: [ericyu: LEVIRCD_Cropped256](#)
- CLCD: [ericyu: CLCD_Cropped_256](#)
- OSCD: [blaz-r: OSCD_RGB_Cropped_96](#)

Data pre-processing details are listed in Appendix E.

B Extended results

B.1 Main results with additional metrics

Table 2 presents the results from the main paper with additional precision and recall metrics. In addition to the discussion in the main body of the paper, we note that ChangeFlow consistently achieves high recall while maintaining a balance between precision and recall. Compared to the previous best overall method, ChangeDINO, our method achieves a recall that is almost 5 percentage points higher. It does, however, achieve lower precision, but when these two are combined in F1, our method achieves a better balance.

Table 2: Change detection results across four diverse datasets and their average. We report Precision (Pr.), Recall (Re.), and F1 score. **First**, **second**, and **third** place results are marked.

	SYSU			LEVIR			CLCD			OSCD			Avg		
	Pr.	Re.	F1	Pr.	Re.	F1	Pr.	Re.	F1	Pr.	Re.	F1	Pr.	Re.	F1
FC-Siam-Diff [14] _{ICIP18}	83.5	61.5	70.8	83.0	80.6	81.8	54.0	54.3	54.1	27.8	68.5	39.4	62.1	66.2	61.5
ChFormer [4] _{IGARSS22}	82.8	73.5	77.9	91.7	87.3	89.5	61.4	60.4	60.8	60.2	40.1	48.1	74.0	65.3	69.1
SwinSUNet [57] _{TGRS22}	89.2	67.2	76.6	86.9	91.7	89.3	79.5	72.5	75.8	61.7	46.3	52.8	79.3	69.4	73.6
GFM [34] _{CVPR23}	89.7	74.3	81.2	90.8	88.8	89.8	82.2	73.2	77.5	55.9	52.5	54.1	79.6	72.2	75.7
GCD-DDPM [55] _{TGRS24}	54.5	78.9	64.5	79.0	82.6	80.7	42.4	52.3	46.9	48.3	3.7	7.0	56.1	54.4	49.8
BiFA [58] _{TGRS24}	87.4	80.4	83.8	90.9	88.1	89.5	79.4	70.1	74.5	61.5	27.1	37.4	79.8	66.4	71.3
MaskCD [56] _{TGRS24}	88.0	80.0	83.8	91.5	89.2	90.3	79.5	73.9	76.6	60.9	24.3	34.7	80.0	66.8	71.4
ChMamba [10] _{TGRS24}	89.6	74.7	81.5	92.4	91.2	91.8	87.3	74.4	80.3	63.4	36.1	45.8	83.2	69.1	74.9
MTP [52] _{JSTARS24}	88.5	75.2	81.3	92.8	90.7	91.7	85.4	75.8	80.3	43.9	66.5	52.8	77.6	77.0	76.5
HySCDG [6] _{CVPR25}	83.3	74.6	78.7	92.6	89.7	91.1	71.1	58.8	64.3	64.4	45.9	53.6	77.8	67.2	71.9
DDPM-CD [3] _{WACV25}	87.3	74.7	80.5	93.1	88.8	90.9	78.9	65.2	71.4	61.9	26.5	37.1	80.3	63.8	70.0
SatDiFuser [22] _{ICCV25}	88.6	76.3	82.0	91.0	89.3	90.2	86.2	73.0	79.1	72.9	44.4	55.2	84.7	70.8	76.6
BTC [44] _{TGRS25}	90.2	75.8	82.4	92.7	90.3	91.5	86.2	76.2	80.9	64.1	47.1	54.3	83.3	72.3	77.3
ChangeDINO [13] _{arXiv25}	88.2	80.1	83.9	92.9	91.4	92.2	86.6	76.7	81.4	73.0	49.3	58.8	85.2	74.4	79.1
ChangeFlow _(10step, 5sep)	86.8	84.3	85.6	91.5	92.7	92.1	86.3	82.8	84.5	62.2	57.0	59.5	81.7	79.2	80.4
ChangeFlow _(1step, 5sep)	86.7	84.5	85.6	91.4	92.6	92.0	86.5	82.6	84.5	65.5	53.7	59.0	82.5	78.4	80.3
ChangeFlow _(10step, 1rep)	88.2	80.0	83.9	92.8	91.3	92.0	87.7	81.3	84.4	65.1	51.5	57.5	83.4	76.0	79.5

B.2 Additional ablations

In this subsection, we present additional ablations. First, we report an additional experiment with a training time-step-sampling alternative. Next, we present different approaches to conditioning vector resizing beyond bicubic interpolation, and finally, we evaluate different normalisation layers beyond LayerNorm. For visual results (including VAE mask reconstructions), refer to Appendix C. Implementation details are in Appendix E.2.

Training time-step sampling approach. ChangeFlow uses logit-normal time sampling during training [17]. This type of sampling emphasises time-steps around 0.5, which is a halfway point between noise and data. This is achieved by sampling from the normal distribution $\mathcal{N}(0, 1)$ and applying a sigmoid to the value, which maps the time to the interval $[0, 1]$. The resulting sampled time is thus concentrated around 0.5, focusing training on the most critical point on the straight line, where paths are most likely to cross and require the most rectification [33]. A commonly used alternative is uniform sampling on the interval $[0, 1]$, which assigns equal probability to all points. For easier visualisation, a histogram of 100,000 sampled steps in both manners is presented in Figure 1

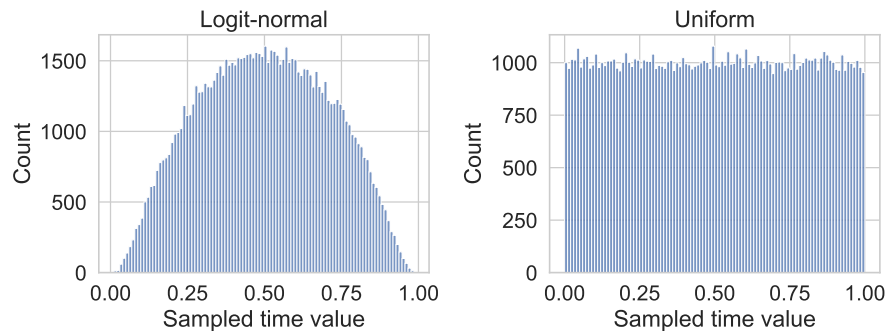


Fig. 1: Histogram of 100,000 sampled timesteps in logit-normal and uniform fashion. ChangeFlow uses logit-normal sampling, which emphasises learning at the critical halfway point between noise and data.

To demonstrate that logit-normal sampling is important for ChangeFlow, we also evaluate a uniform alternative and present the results in Table 3. Uniform sampling consistently performs worse across all datasets, underscoring the importance of focusing training on the more critical halfway point in the rectified flow.

Conditioning resizing. Since the spatial dimension of the VAE latent space may not align with that of the image encoder, some form of resizing is required.

Table 3: Training time-step sampling ablation results.

	SYSU	LEVIR	CLCD	OSCD	Avg.
	F1	F1	F1	F1	F1
Logit-normal t sampling (<i>Ours</i>)	85.6	92.1	84.5	59.5	80.4
Uniform t sampling	85.5	91.5	83.5	57.8	79.6

In our case, the height and width dimensions of the DINOv3 encoder latent are half the size of the VAE’s latent (downsampling by 16 vs 8). To match the dimensions, we use bicubic interpolation to rescale the conditioning vector (which comes from features from the image encoder). We also explored some alternatives, with results presented in Table 4. A future possibility would also be some form of learnable upscaling. Current results indicate that bicubic achieves the best overall performance, while bilinear outperforms it on SYSU. Lanczos performs worst overall, but all 3 approaches are relatively similar, indicating that this choice is important but not to the extent of other architectural decisions, such as normalisation layers.

Table 4: Conditioning vector resizing approach ablation.

	SYSU	LEVIR	CLCD	OSCD	Avg.
	F1	F1	F1	F1	F1
Bicubic (<i>Ours</i>)	85.6	92.1	84.5	59.5	80.4
Bilinear	85.9	92.1	83.9	57.8	79.9
Lanczos	84.5	92.1	83.8	58.8	79.8

Different normalisation layers. ChangeFlow uses LayerNorm [1] for conditioning feature normalisation, a common normalisation layer in recent architectures. We also evaluated two other options: InstanceNorm and BatchNorm. Results are presented in Table 5. LayerNorm achieves the best overall performance, while both alternatives perform considerably worse. This can be explained by the general properties of normalisation layers: BatchNorm depends on batch statistics and can introduce instability when the batch contains heterogeneous bi-temporal pairs, whereas InstanceNorm removes instance-specific contrast information useful for change detection. In contrast, LayerNorm normalises features along the channel dimension of each spatial location independently of other samples, preserving per-pixel structure while ensuring consistent feature scaling. These properties make LayerNorm particularly well-suited for conditioning generative models, yielding the strongest performance in our setting.

Table 5: Normalisation layer ablation.

	SYSU	LEVIR	CLCD	OSCD	<i>Avg.</i>
	F1	F1	F1	F1	F1
LayerNorm (<i>Ours</i>)	85.6	92.1	84.5	59.5	80.4
InstanceNorm	84.0	91.6	80.8	56.6	78.3
BatchNorm	84.4	92.0	83.7	56.2	79.1

B.3 Additional confidence results

In this subsection, we present additional analyses of the confidence our model produces through ensemble aggregation. First, we analyse the rule for binarising/thresholding predictions in the ensemble based on prediction agreement (2-predictions-equal-change). We then present additional qualitative results, followed by the quantitative evaluation. Implementation details for our analyses are in Appendix E.2.

Optimal confidence threshold. To find the optimal threshold for binarising our predicted ensemble of change masks, we evaluate different thresholds on **the validation** set. Since OSCD does not contain a validation set, we skip it. *The optimal binarisation regime in an ensemble of five predictions is to consider a region changed if at least two predictions mark it as such*, as shown in Figure 2. Figure 3 illustrates that this point represents the best precision-recall trade-off, but the model offers the option to either prefer recall or precision by varying this threshold. For fair evaluation, we use the optimal technique in terms of F1 (2-predictions-equal-change) from the validation set to evaluate ChangeFlow on the test set in the main paper.

Additional qualitative results. Additional confidence visualisations with accompanying discussion are in Appendix C.5.

Quantitative evaluation. Besides the above analysis, we also try to quantitatively assess the usefulness of this confidence signal. We compute *error-AUROC*, which measures how well low-confidence pixels predict actual binary errors. To compute this metric, we treat each pixel’s *error indicator* (1 if the predicted label differs from the ground truth, 0 otherwise) as the binary target, and the *confidence score* derived from sampling agreement as the predictor. We convert confidence into an *error score* by taking 1 minus its value, so that low-confidence pixels correspond to higher predicted error.

Table 6 reports dataset-level results. ChangeFlow achieves relatively high error-AUROC values ranging from 0.70 to 0.87, indicating that its confidence identifies error-prone pixels in most cases. The performance is lowest on SYSU, where the regions are large, so many pixels can be misclassified if the region-level decision is incorrect.




Fig. 2: Different thresholds used for binarisation evaluated on **validation set**. OSCD does not contain a validation set, so we skip it. The best performance is achieved by binarising all regions where at least two ensemble predictions indicate a change.




Fig. 3: Different thresholds used for binarisation evaluated on **validation set** (3 datasets average) in terms of precision and recall. OSCD does not contain a validation set, so we skip it. The selection of the binarisation threshold offers a trade-off between recall and precision. The optimal precision-recall trade-off occurs at the 2-predictions-equal-change rule.

Table 6: Dataset-level error-AUROC for ChangeFlow. Error-AUROC measures how well low-confidence pixels (derived from sampling-based agreement) predict true errors. Higher is better.

Dataset	Err-AUROC \uparrow
OSCD96	0.870
LEVIR	0.864
CLCD	0.769
SYSU	0.703

The confidence mechanism is particularly beneficial for downstream applications requiring conservative or safety-aware predictions. Overall, these results demonstrate that ChangeFlow not only improves raw F1 accuracy but also provides a meaningful confidence signal that can be leveraged for selective prediction or human-AI collaboration workflows. While this confidence is not ideal and offers many opportunities for further improvement, it has many benefits that prior discriminative methods do not offer, namely the precision-recall trade-off (Appendix B.3), benefits from ensembling, and more expressive visual feedback (Appendix C.5).

C Additional qualitative results

This section provides additional qualitative results. First, we present the extended main qualitative results in comparison with a wider selection of related work (Appendix C.1), followed by some failure cases (Appendix C.2). Next, we

present visual VAE mask reconstructions (Appendix C.3) and visualisations of intermediate generation steps (Appendix C.4). Finally, we present additional confidence visualisations (Appendix C.5).

C.1 Main qualitative results

Figure 4 presents visual results in comparison to a wider set of related methods. ChangeFlow excels at predicting more coherent change masks and capturing full changed regions (low number of false negatives). No prior method can consistently match this behaviour across multiple datasets, as also reflected in ChangeFlow’s superior recall (see Appendix B.1).

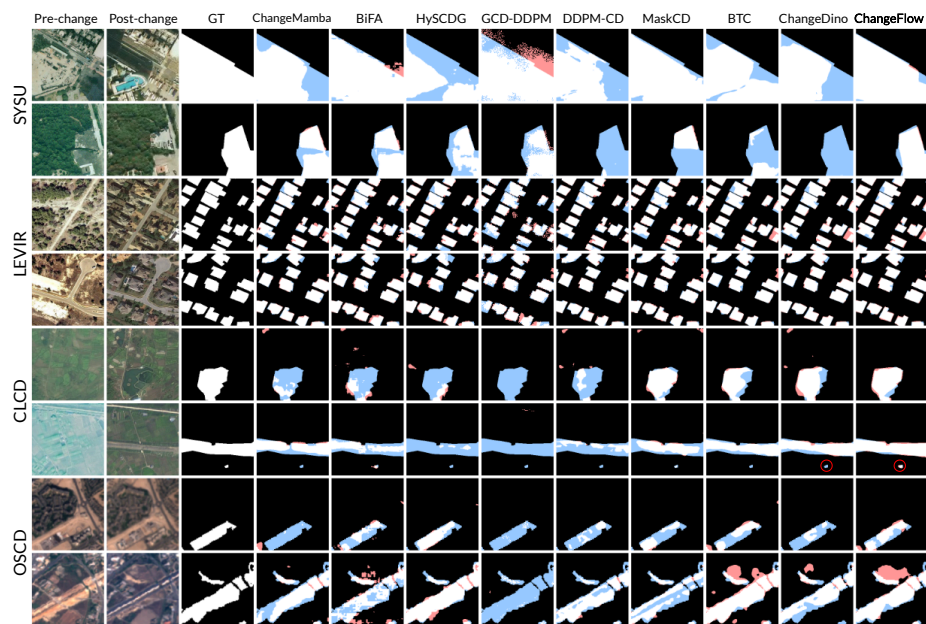


Fig. 4: Additional qualitative results.

C.2 Failure cases

Figure 5 contains a visualisation of some failure cases. ChangeFlow does miss some changed regions in specific situations, but the visualisations show that most other methods struggle with similar problems. The hardest example is shown in the CLCD row, where no model correctly predicts the majority of the changed region, indicating its high semantic nature and difficulty. In the first two rows (SYSU and LEVIR), we see that some of these changes may be due to mislabelling. Later (see additional confidence visualisations below), we show that the model is quite uncertain about the misclassified LEVIR case.

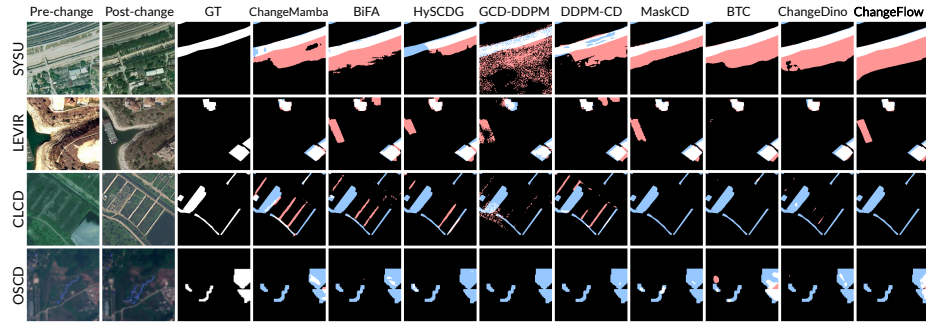


Fig. 5: Additional failure qualitative results.

C.3 Visualisation of VAE mask reconstruction

Our method uses a pretrained variational autoencoder (VAE) from SD-XL [41]. This network was originally trained on RGB images, so it is immediately obvious that we can also encode binary change masks with minimal loss of data. We verified this and presented the results in the main paper (Section 4.1), with minimal drop in F1 score and mean absolute error. In Figure 6, we also support the quantitative results with visual proof that the VAEs trained for RGB images sufficiently encode binary masks. The implementation details for the analysis are in Appendix E.2. Even if this whole observation might not be intuitive at first glance, we note that a binary image is a special case of an RGB image, so it is natural that a certain subspace of the VAE latent space encodes such imagery.




Fig. 6: Original (top row) and reconstruction (bottom row) of binary change masks through pretrained SD-XL VAE. The encoding process preserves the details and structure of masks with minimal data loss. Refer to the main paper Section 4.1 for quantitative evaluation.

C.4 Intermediate steps visualisation

Figure 7 presents a visualisation of all 10 generation steps to complement the 5 with a stride of 2 in the main paper. As we can see, and as we have quantitatively

evaluated in the main paper, the coherent region appears quite early in the process. The final steps then predominantly focus on border refinement (we recommend zooming in to see this clearly).




Fig. 7: Visualisation of intermediate steps in the latent generative mask prediction. ChangeFlow iteratively predicts from pure noise to a binary mask. Here, we decode the intermediate latent representation into a binary mask at each step.

C.5 Additional confidence visualisations

Figure 8 presents 4 additional cases of confidence-based prediction visualised. This effect is especially pronounced in the case of SYSU and CLCD. Multiple predictions in the ensemble ensure that all changed regions get coherently predicted. In the case of LEVIR (an example from a previous failure case), we can see that the building-like object on the lake (which the annotators did not consider a change) yields a relatively uncertain prediction (low intensity) and could have been avoided with stricter binarisation conditions.

D Computational efficiency

Details of our computational-efficiency evaluation protocol are provided in Appendix D.1. We additionally report complementary metrics for our method and selected state-of-the-art baselines, including GFLOPs and inference time in Appendix D.2.

D.1 Protocol

We report three efficiency metrics: parameter count, inference time (also expressed as frames per second, FPS), and GFLOPs. GFLOPs are measured with the official PyTorch profiler.

The protocol closely follows the one from [44]. Inference time is measured using a pair of 256×256 RGB inputs. All models, except ours, are evaluated

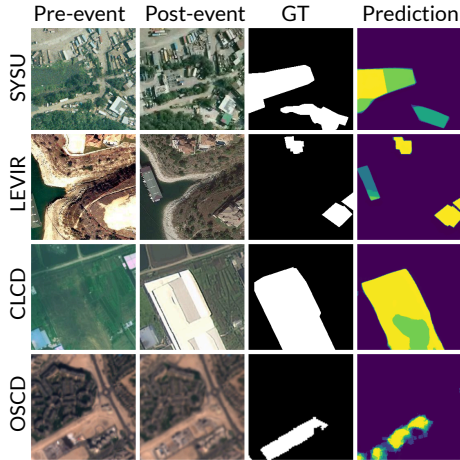


Fig. 8: Additional visualisations of ensembled predictions where prediction agreement yields confidence with respect to change class.

in `float16` where supported. Our model does not support `float16` for all modules; therefore, we use `torch compile` when measuring inference time. To robustly measure the metrics, we perform 1000 warm-up forward passes followed by 1000 timed forward passes; this procedure is repeated five times, and we report the average runtime per forward pass. All measurements are conducted on an NVIDIA A100-SXM4 40GB GPU (and AMD Epyc 7H12 CPU) on a Slurm cluster. We use the same protocol for all methods and efficiency results.

D.2 Extended computational results

Table 7 shows extended computational efficiency results including frames per second (FPS) derived from inference time, parameter count, floating operations per second (FLOPS), and change detection metrics. ChangeFlow outperforms ChangeDINO by 1.3 points in change detection but offers comparable throughput and inference time. Some other discriminative methods, such as BTC and HySCDG, do achieve faster inference, but they do not perform nearly as well as ChangeFlow in change detection.

Due to multiple inferences, ChangeFlow has a higher FLOPS value, in line with related diffusion-based models (DDPM-CD and SatDiFuser). *This is not an issue on modern hardware with massive parallelisation capabilities*, as reflected by FPS metric. We argue that FPS is a much better measure of the model’s actual inference speed, as reflected in the comparison between our `1step5rep` and `10step1rep` setups: the FLOPS value is higher (due to parallelisation) for the multiple-repetition setup, but inference is actually faster, as these repetitions can be easily parallelised on modern GPUs.

The previous generative method, GCD-DDPM, is almost 3 orders of magnitude slower than our method. This stems from its complex conditioning scheme,

which uses the auxiliary CD architecture’s output for guidance. It also operates in pixel space and does 1000 generation steps, as opposed to 10 latent steps in ChangeFlow.

Another observation is that, even though ChangeFlow is a true generative method, it operates faster than DDPM-CD and SatDiFuser, which use a generative (diffusion) network solely as a feature extractor, while achieving substantially higher change detection accuracy.

Table 7: Computational efficiency results for each model. We report FPS (derived from inference time), inference time, parameter count, GFLOPs, and average Precision, Recall, and F1 across 4 datasets. All results were obtained using an Nvidia A100-SXM4 40GB GPU using the above-described protocol.

	FPS	Inference Time	Param.	FLOPS	<i>Avg</i>		
	[img/s]	[ms]	[M]	[10 ⁹]	Pr.	Re.	F1
FC-Siam-Diff [14]	170.1±0.0	5.9±0.0	1.4	4.6	62.1	66.2	61.5
ChFormer [4]	36.2±0.0	27.6±0.0	41.0	234.6	74.0	65.3	69.1
SwinSUNet [57]	33.1±0.1	30.2±0.1	43.6	32.6	79.3	69.4	73.6
GFM [34]	44.9±0.1	22.3±0.1	120.5	109.2	79.6	72.2	75.7
GCD-DDPM [55]	0.02±0.0	43563.6±0.0	131.9	531997.7	56.1	54.4	49.8
BiFA [58]	32.2±0.0	31.0±0.0	9.9	4.3	79.8	66.4	71.3
MaskCD [56]	6.5±0.0	153.5±1.0	107.4	143.2	80.0	66.8	71.4
ChMamba [10]	14.4±0.0	69.6±0.2	92.4	96.2	83.2	69.1	74.9
MTP [52]	31.2±0.0	32.1±0.0	107.8	196.9	77.6	77.0	76.5
HySCDG [6]	41.0±0.1	24.4±0.0	65.1	64.8	77.8	67.2	71.9
DDPM-CD [3]	4.6±0.0	217.6±0.4	437.5	8871.2	80.3	63.8	70.0
SatDiFuser [22]	1.8±0.0	542.2±0.9	1413.6	6142.9	84.7	70.8	76.6
BTC [44]	32.4±0.0	30.8±0.0	120.1	221.4	83.3	72.3	77.3
ChangeDINO [13]	8.9±0.5	112.3±6.2	311.1	1269.1	85.2	74.4	79.1
ChangeFlow (10step, 5rep)	8.1±0.0	124.2±0.0	403.3	4673.9	81.7	79.2	80.4
ChangeFlow (1step, 5rep)	18.2±0.0	55.0±0.0	403.3	3543.1	82.5	78.4	80.3
ChangeFlow (10step, 1rep)	15.3±0.1	65.2±0.5	403.3	1188.4	83.4	76.0	79.5

E Extended implementation details

This section contains detailed implementation details for our model in Appendix E.1, implementation details for our ablations and analysis (coherence) in Appendix E.2, and finally also details regarding the related methods in Appendix E.3.

All experiments, including all our experiments and related method execution, were conducted on an NVIDIA A100-SXM4 40GB GPU (and AMD Epyc 7H12 CPU) on a Slurm cluster.

E.1 Our model - ChangeFlow

This subsection contains additional implementation details not included in the main paper for the modules used in ChangeFlow: the SD-XL VAE, the diffusion transformer (DiT), the DINOv3 image encoder, feature difference and normalisation, ensembling details and other training-related details.

SD-XL VAE. We use the VAE from SD-XL (Stable Diffusion XL) [41] for image generation. It is kept frozen in the base model, so no gradient flows through the encoder or the decoder. We selected it for its compact 4-channel latent space ($d = 4$). We also ablated this choice as useful in ablation studies of the main paper. Specifically, we use the HuggingFace [stabilityai/sdxl-vae](#) version and keep all details unchanged. We utilise the scaling factor and, as standard [41], multiply the latent by 0.13025 and then, before decoding, divide by it. As already explained in the main paper, to convert a binary mask to RGB, we simply repeat the value along the channel dimension. When decoding, we simply average the 3 RGB channels to retrieve a single-channelled binary mask. Since the VAE expects images to be normalised to the range $[-1, 1]$, we rescale all masks to this range before encoding, then back to $[0, 1]$ after decoding.

DiT model. As the model that predicted the velocity field, we opt for the recent diffusion transformer (DiT). The architecture itself is based on [LLaMA-2 DiT](#), with the implementation adopted from the [minRF](#) GitHub repo. We set the channel dimension to 256, use 10 layers with 8 heads each, and a patch size of 1. We do not use class embeddings or classifier-free guidance. The input channel dimension is set to the sum of the image encoder dimension c and the VAE latent dimension d , specifically $1024 + 4$, for a total of 1028, since the model receives a concatenation of feature difference and noise in the shape of a mask VAE latent (see Section 4.1). Output channel dimension is set to VAE latent dimension (4, which matches the latent of the expected output mask latent). The model also takes the time-step value in the range $[0, 1]$ as an input, which is then embedded using the TimeEmbedder (see the repo mentioned above for details). Other DiT hyperparameters remain the same as in the repo mentioned above. This module has an initial learning rate of $1 \cdot 10^{-4}$.

Image encoder. We use DINOv3 ViT-L as a shared weight image encoder, specifically the version from HuggingFace [facebook/dinov3-vitl16-pretrain-lvd-1689m](#). It has a hidden dimension of 1024, a patch size of 16, and 24 layers. We do not modify any default hyperparameters and use the provided image normalisation parameters. The model is finetuned during training following [44], and we set the learning rate of this module to $5 \cdot 10^{-5}$. We extract the features from the last (24th) layer. We discard register and class tokens and reshape features from $\mathbb{R}^{l \times c}$ to $\mathbb{R}^{h' \times w' \times c}$, where $h' = w' = \sqrt{l}$ and $c = 1024$.

Feature difference and normalisation. To obtain a conditioning vector, features extracted with the above-described image encoder are normalised before subtraction (differencing). We opt for LayerNorm [1], a standard choice and the best performer according to ablations in Appendix B.2. This is applied across channel (embedding) dimension c , in the feature map $\mathbb{R}^{h' \times w' \times c}$. LayerNorm hyperparameters remain default as in PyTorch, and the trainable scale parameters have the same learning rate as DiT: $1 \cdot 10^{-4}$.

Feature difference is computed per element, meaning that given two feature maps, both of shape $\mathbb{R}^{h' \times w' \times c}$, we subtract the values at the same indices of h', w', c . In our base model, we then apply the absolute value to this difference. This absolute difference represents our conditioning vector.

Finally, the conditioning vector is resized to match the VAE spatial dimensions (in our case, we upscale h' and w' by a factor of 2) using simple bicubic interpolation (PyTorch implementation). This choice of resizing method is ablated in Appendix B.2.

Ensembling details. To obtain an ensemble of predictions, we repeat the inference N times, where $N = 5$ in our base model, the same way as explained for a single inference. More specifically, this means that we sample N noise vectors $\mathbb{R}^{h \times w \times d} \sim \mathcal{N}(0, 1)$ that have the same shape as the expected mask VAE latent. These then undergo the standard 10-step inference via ODE integration, resulting in N final change-mask latents: $\{\hat{x}_i | i \in 0..N\}$. These are then individually decoded via the VAE decoder, and the RGB channels are averaged to obtain single-channel binary masks, resulting in a final ensemble of binary change masks $\{\hat{M}_i | i \in 0..N\}$. We then stack predictions in a new dimension to obtain $\hat{M}_{ens} \in \mathbb{R}^{N \times h \times w}$ and aggregate via averaging across the new dimension to obtain a final prediction $\hat{M} \in \mathbb{R}^{h \times w}$.

During inference, this process can be easily parallelised since the *repetitions are independent*. By stacking the N different initial noise vectors $\{x_o^i | i \in 0..N\}$ in a new dimension to get $x_o^{batched} \in \mathbb{R}^{N \times h \times w \times d}$, the ODE integration is performed in batched manner, resulting in a batched final latent $\hat{x}^{batched} \in \mathbb{R}^{N \times h \times w \times d}$, which is then decoded and merged as explained above. This means that increasing repetitions increases inference time with respect to the parallelisation capabilities of modern hardware, in theory enabling a smaller overhead with better parallelisation.

Other details. As already explained in the main paper and above, we obtain a single channel prediction from the ensemble of N predictions $\mathbb{R}^{N \times h \times w}$ by averaging across the ensemble dimension N to get $\mathbb{R}^{h \times w}$. The values in this prediction are continuous but represent 5 different hypotheses. To achieve the effect of two predictions indicating a change in continuous space, we set the threshold to 0.3. This is equivalent to discretising into 5 values (with rounding) and then thresholding at ≥ 2 . This was established as the optimal threshold on the validation set with results presented in Appendix B.3.

As explained in the main paper, we sample time steps during training in a logit-normal fashion (see Appendix B.2 for details and ablations). During inference, time steps are equally spaced on the interval $[0, 1]$. In our case, we use 10 timesteps ($T = 10$). The value was simply selected as the one where performance according to ablations is good, but the same ablation also shows that increasing T does not yield consistent gains. The number of repetitions in the ensemble (i.e., 5) was selected as it represents a good speed-performance trade-off. While we could’ve selected a higher value to achieve even better CD performance, we believe our selection is a fair choice given its similar inference speed to the previous best method.

As already explained in the main paper, we use rotation and flipping augmentations, each applied with a probability of 30%. All input images are of size 256×256 , which means that for OSCD, we rescale the images from crops of 96×96 , following other works [44, 52]. Data normalisation is specified above for the image encoder and VAE. Dataset details are in Appendix A.

A cosine scheduler without restarts is used in all cases, with the PyTorch default implementation. The Muon optimiser comes from the Timm library. We picked this option with the recent success of LLM applications, but the change in results compared to AdamW was minimal in preliminary studies.

Metric implementations come from TorchMetrics and augmentations from Albumentations.

E.2 Ablation and analyses implementation details

Encoder ablation details. All parameters stay the same as for the base model, except for the following, which are specific to encoder selection. For DINOv2, we use [facebook/dinov2-large](#); all hyperparameters stay the same. For the DINOv3 satellite, we use [facebook/dinov3-vitl16-pretrain-sat493m](#), all hyperparameters stay the same. For RADIO 2.5, we use [nvidia/RADIO-L](#); all hyperparameters stay the same, except the learning rate, which is divided by 10, and normalisation is set to author-provided. For RADIO 4, we use [nvidia/C-RADIOv4-SO400M](#) (ShapeOptimised version since there is no ViT-L), and the hyperparameters are the same as in RADIO 2.5.

Conditioning ablation details. The process of base-feature normalisation and differencing is explained in the implementation details above for ChangeFlow. For other ablated options, we list the details here.

Feature difference is computed per element, meaning that given two feature maps, both of shape $\mathbb{R}^{h' \times w' \times c}$, we subtract the values at the same indices of h', w', c . When we compute a signed difference, we subtract the feature map of the second image (the one at a later time step) from that of the first. In the case of concatenation conditioning vector, we concatenate the features in the channel dimension to obtain $\mathbb{R}^{h' \times w' \times 2c}$ (and accordingly adjust the DiT input channel dimension).

In the case of L2 normalisation, we compute the L2 vector norm across the channel dimension of feature map $\mathbb{R}^{h' \times w' \times c}$ (resulting in $\mathbb{R}^{h' \times w'}$ norm vector),

then divide all corresponding channel values by this norm. Unlike LayerNorm, this option does not contain the learnable scale parameters.

VAE ablation details. We use the following VAEs from Huggingface and leave all hyperparameters the same as the original: [stabilityai/sd-xl-vae](#), [stabilityai/stable-diffusion-3.5-medium](#), [black-forest-labs/FLUX.1-dev](#), and [Tongyi-MAI/Z-Image-Turbo](#). The encoder is always frozen, while the decoder is frozen except in ablations where indicated. Since the input of DiT is defined as the sum of the VAE latent dimension d and the image encoder latent dimension c , the VAE part of the dimension is accordingly changed to the latent dimension of VAE: $d = 4$ for SD-XL and $d = 16$ for all others.

In experiments where we also finetune the SD-XL VAE decoder, we use a standard binary dice loss (same as in [44]) on the change mask, computed with single-step single-repeat inference and binarised from RGB to a single channel. The gradient passes through both the VAE decoder and the DiT, enabling us to avoid the standard rectified flow MSE loss in the "Pixel loss only" experiments. All other parts of DiT and the image encoder keep the same configuration in these experiments. We set the VAE decoder learning rate to the same as DiT's: $1 \cdot 10^{-4}$. In the case of the CNN decoder, it is a UNet-like model with a single final CNN block with a channel dimension of 256. Its weights are randomly initialised. Even with the CNN decoder, we keep the VAE encoder for target mask encoding. The learning rate, gradient propagation, and loss are the same as in the finetuned VAE case explained above.

Coherence analysis details. In the main paper, we perform two analyses: one to calculate the deviation from the expected number of holes, and the other to calculate the deviation from the expected connected component (CC) count. We use these metrics to qualitatively evaluate coherence based on the fact that coherent prediction should: (i) match the expected changed region count (CC deviation metric) and (ii) not contain sporadic holes in change masks (hole deviation metric).

To compute these metrics, we operate on the binary change mask produced by each method. For the *CC metric*, we extract all foreground connected components using the SciPy library (default parameters) and retain only those with an area exceeding a small threshold (10 pixels in our experiments). Components below this threshold are considered spurious fragments and removed. We then count the remaining components and compare this count to the ground-truth CC count (obtained the same way but with ground truth mask), yielding the connected-component deviation.

For the *hole metric*, we apply the same procedure to the *background*: we identify all background connected components and discard those that touch the image border, since such regions represent true background rather than holes. Among the remaining enclosed components, we keep only those whose area exceeds the same minimum threshold, and their count forms the hole count for that prediction. The difference between this value and the ground-truth hole

count (obtained as explained above but with ground truth mask) yields the hole deviation.

Both metrics quantify structural coherence by penalising either unnecessary fragmentation (extra CCs) or unwanted perforation of change regions (holes). Lower deviation indicates that a method produces more globally consistent change masks.

For the experiment where we calculate these metrics at the intermediate steps of the ChangeFlow masks, we use the exact same approach as above, applied to masks decoded from intermediate latents and thresholded as in the base model.

E.3 Related methods implementation details

For all models, we use the same data as in our case. We do adopt the normalisation and other model-specific settings for data processing.

We use the official code, hyperparameters, and weights provided by the authors for all evaluated remote sensing foundation models. The specific versions of the code used are as follows (repo + commit):

- GFM [34]: [GFM commit: 4dd248e8544b3b6a49f5173b0931d97a17a7f424](#)
- MTP [52]: [MTP commit: 962f7fd8781c095eb26db65ead3016e666b6d417](#)
- SatDiFuser [22]: [MTP commit: 962f7fd8781c095eb26db65ead3016e666b6d417](#)

Since foundation models lack a predefined, exact change-detection architecture, we adopt the authors’ architecture code and load the weights as the *encoder* into the BTC framework [44]. The configuration for MTP and GFM is the same as in [44]. For SatDiFuser, we use the default parameters and UPerNet decoder with simple feature difference, similar to BTC [44].

We use official code, hyperparameters, and weights (where applicable) for all change detection methods. The following are repos and commits:

- FCS-Diff [14]: [fully_convolutional_change_detection commit: 4dd83231f25-319a7ebb16cbfa9912541ceabac9a](#)
- ChangeFormer [4]: [ChangeFormer commit: afd1b7ed640aa265a2c730de9584-16ae7356a2f9](#)
- SwinSUNet [57]: [SwinSUNet commit: 721daf84238eda40fb49d626c21df4ed-2246aa9e](#)
- GCD-DDPM [55]: [GCD commit: ecf2f25c55e849dc92d948e6ed0ed9ff05163b-96](#)
- BiFA [58]: [BiFA commit: 56cd0da461e5e4b0d6a9b4f3321f0a81a91d21b8](#)
- MaskCD [56]: [MaskCD commit: 31e3e15c50a81a369fc7fec2134b61fbedaa60-05](#)
- ChangeMamba [10]: [ChangeMamba commit: a91b82ee45059ce159f5f6f5d8e-5818c33b84e68](#)
- HySCDG [6]: [HySCDG commit: 05db2154dc9f24ee650fb27285617abbf38d-8a9e](#), pretrained model weights from HF: [Yanis236/FSC-Pretrained](#)

- DDPM-CD [37]: [ddpm-cd commit: 4970792f65227958ffaa1de787649ce2c58-39f12](#)
- BTC [44]: [BTC-change-detection commit: db41090f2f26b84b2bd803b5177-56a10f0805b2f](#)
- ChangeDino [13]: [ChangeDINO commit: 1870b2641b0eb83d367a484e53e023-b578b26c1f](#)

We keep all hyperparameters the same as those set by the authors, except for the epoch count on SYSU and OSCD, where we perform some tuning to improve performance given the dataset size differences.