


Full Length Article

Continual learning of conjugated visual representations through higher-order motion flows

Simone Marullo ^{a,c}, Matteo Tiezzi ^b, Marco Gori ^c, Stefano Melacci ^c *

^a Department of Information Engineering, University of Florence, Via di S. Marta 3, Florence, 50139, Italy

^b Istituto Italiano di Tecnologia, Via Enrico Melen 83, Genova, 16163, Italy

^c Department of Information Engineering and Mathematics, University of Siena, Via Roma 56, Siena, 53100, Italy

ARTICLE INFO

Keywords:

Learning from constraints
Lifelong learning
Feature extraction from a single video stream
Online learning

ABSTRACT

Learning with neural networks from a continuous stream of visual information presents several challenges due to the non-i.i.d. nature of the data. However, it also offers novel opportunities to develop representations that are consistent with the information flow. In this paper we investigate the case of unsupervised continual learning of pixel-wise features subject to multiple motion-induced constraints, therefore named *motion-conjugated feature representations*. Differently from existing approaches, motion is not a given signal (either ground-truth or estimated by external modules), but is the outcome of a progressive and autonomous learning process, occurring at various levels of the feature hierarchy. Multiple motion flows are estimated with neural networks and characterized by different levels of abstractions, spanning from traditional optical flow to other latent signals originating from higher-level features, hence called higher-order motions. Continuously learning to develop consistent multi-order flows and representations is prone to trivial solutions, which we counteract by introducing a self-supervised contrastive loss, spatially-aware and based on flow-induced similarity. We assess our model on photorealistic synthetic streams and real-world videos, comparing to pre-trained state-of-the-art feature extractors (also based on Transformers) and to recent unsupervised learning models, significantly outperforming these alternatives.

1. Introduction

Learning from a continuous stream of visual data is a longstanding challenge for the machine learning community and AI in general. The challenging nature of this task becomes evident when we consider an artificial agent which continuously processes the streamed data, learning in an online manner (Tiezzi, Marullo, Faggi et al., 2022). The recent outbreak of self-supervised techniques in computer vision has closed the gap with supervised baselines in many vision tasks (He, Fan, Wu, Xie, & Girshick, 2020), leveraging the intuition of relating different views of the same entity, usually with the goal of building image-level representations, but without explicitly relying on motion. Most of the existing approaches are based on contrasting the similarity of positive examples with the dissimilarity of negative examples, using memory banks or large batches (Chen, Kornblith, Norouzi and Hinton, 2020). While contrastive learning works surprisingly well for a variety of downstream tasks (Le-Khac, Healy, & Smeaton, 2020), typical models require offline training on very large collections of images and some kind of prior knowledge.

Our starting point is the observation that unlike machines, humans and animals do not learn by going through vast and randomly shuffled sets of images. Instead, they learn continuously, as they experience the world around them, without storing and labeling every piece of information they come across. In this paper, we argue that machines too can potentially adopt this natural way of learning, acquiring knowledge continuously from their surroundings, without depending on large pre-existing datasets. Although this goal might seem distant, this work is a small step in that direction.

The first central point we consider in this paper is the opportunity of substantially relying on motion to design a natural contrastive framework for agents that progressively learn from a visual stream. Foundational studies in vision and perception have shown (Spelke, 1990) that the presence of motion significantly enhances the ability of biological perceptual systems to identify and segment visual patterns into specific entities. Empirical evidence (Ostrovsky, Meyers, Ganesh, Mathur, & Sinha, 2009) supports the idea that the capability of parsing static visual scene is systematically achieved much later than the one of parsing dynamic scenes. As a matter of fact, the Gestalt Principles

* Corresponding author.

E-mail address: stefano.melacci@unisi.it (S. Melacci).

of common fate (Wertheimer, 1938) hypothesized the role of motion as a fundamental cue for visual perception in the early 20th century. Such intuition has been exploited for designing simple pretext tasks in the context of unsupervised learning (Mahendran, Thewlis, & Vedaldi, 2019), aligning the similarity between pairs of feature vectors to the similarity between the corresponding flow vectors. The authors of Pathak, Girshick, Dollár, Darrell, and Hariharan (2017) used segments from low-level motion-based grouping to train convolutional networks, yielding easily transferable representations. The model proposed in Tiezzi, Marullo, Faggi et al. (2022) learns from a video stream, exploiting motion and focus-of-attention provided as external cues.

In order to develop the aforementioned visual agents, the temporal dimension is of utmost importance, that is the second focal aspect of this paper. Recently, a lot of emphasis has been put on neural models able to learn over time (Tiezzi, Marullo, Betti et al., 2022; Tiezzi, Melacci, Betti, Maggini, & Gori, 2020), in the attempt of overcoming the conventional i.i.d. assumption and offline learning (i.e., all training data are instantly available and sampled from a static distribution in an independent manner). Most of the attention has been focused on continual supervised learning, with few notable unsupervised exceptions (Betti, Faggi et al., 2022; Madaan, Yoon, Li, Liu, & Hwang, 2022; Rao et al., 2019). Multiple families of approaches have emerged (van de Ven, Tuytelaars, & Tolias, 2022), for which the reader can find a comprehensive description in Lange et al. (2021). Despite the large variety of proposals, learning over time in a continual manner is still a very challenging learning setting, especially when not pushing the emphasis on the memorization of large buffers of past experiences. Regularization techniques are indeed very powerful, and motion seems to offer a natural way of regularizing over time.

Motivated by the aforementioned considerations, we propose to face the learning problem with a neural architecture designed to jointly learn to estimate motion and extract motion-conjugated features (representations) (Betti, Gori and Melacci, 2022) in a continual manner, which we refer to as Continual MOTion-based Self-supervised Feature ExTractor (CMOSFET). Motion can be the outcome of virtually tracking low-level representations (pixel intensities) or higher level features computed by the deep architecture, thus we go beyond the specific notion of optical flow and we introduce higher-order motion flows. We build on the experience of Tiezzi, Marullo, Faggi et al. (2022) and we propose a more sophisticated self-supervised motion-driven contrastive criterion over time and space, learning to estimate multiple motion fields from scratch, without any additional external inputs. In the proposed framework, not only optical flow is estimated in a continual learning fashion, as in Marullo et al. (2022), but also higher-order flows are continuously estimated, computed according to features extracted at several levels in the feature hierarchy. The idea of exploiting multiple motion cues to guide or support learning is the subject of recent theoretical studies (Betti, Boccignone, Faggi, Gori, & Melacci, 2021; Betti, Gori et al., 2022).

The main contributions of this paper, which is an extension of our previous manuscript (Tiezzi, Marullo, Betti, Casoni and Melacci, 2024), are the following: (i) we propose CMOSFET, a model which processes frames in an online manner with a two-branch neural architecture, continuously-and-jointly learning to extract pixel-wise motion fields and pixel-wise visual features, in a self-supervised manner and at multiple levels of abstraction (differently from Tiezzi, Marullo, Betti et al. (2024), where a single level is considered). (ii) We propose a contrastive criterion which revisits the definition of positive and negative pairs, establishing spatio-temporal correspondences/contrasts among pixels belonging to consecutive frames. We further propose a mechanism tailored to tightly couple the development of features and motion estimation with cross-layer relations. (iii) We make learning affordable by restricting the contrastive loss to a group of stochastically selected pixels, proposing to bias the sampling procedure by the way motion and features are predicted over the frame area. (iv) We exploit Motion coherence to naturally regularize feature learning over time,

and we gain further stability and reduce catastrophic forgetting by introducing a fast learner/slow learner implementation, where the slow learner model smoothly adapts to the video stream via a momentum-based update scheme. (v) Our experiments on a recently introduced benchmark and real-world videos confirm the validity of CMOSFET, and the usefulness of higher-order motion flows. CMOSFET outperforms the most related models and shows competitive results when compared to convolutional and Transformer-based networks, even if pre-trained on large collections of images. We experimentally compare our previous proposal (Tiezzi, Marullo, Betti et al., 2024), hereinafter referred to as CMOSFET-Single, with the CMOSFET approach presented in this paper and demonstrate that the latter shows significant improvements.

Our work is organized as follows. Section 2 draws connections with related works. Section 3 describes the proposed model. Different features of what we propose are described in multiple subsections within Section 3. Experiments and the results we obtained are discussed in Section 4, jointly with the main limitations of what we propose. Finally, conclusions are drawn in Section 5, with suggestions for future work.

2. Related work

In the following, we briefly discuss the positioning of the proposed method with respect to related literature domains.

Continual Learning. Recently, a lot of emphasis has been put on neural models able to learn over time (Tiezzi, Marullo, Becattini and Melacci, 2024), with the aim of eliminating the conventional i.i.d. assumption and offline learning (i.e., all training data is available from the start and sampled from a static distribution in an independent manner). Most of the attention has been focused on continual supervised learning, with few notable unsupervised exceptions (Madaan et al., 2022; Rao et al., 2019). Multiple families of approaches have emerged (van de Ven et al., 2022), namely context-specific components, parameter regularization, functional regularization, replay-based methods. The goal of CMOSFET is to learn to predict class labels for pixels in a video stream without a traditional offline training procedure. Such goal is not fully defined in advance, since the agent becomes aware of classes as the human supervisor tells him, and it does not go through a sequence of distinct tasks with clear boundaries—being close to CL task-free scenario (Aljundi, Kelchtermans, & Tuytelaars, 2019). Moreover, the spatiotemporal loss functions we propose introduce a motion-induced regularization effect, arguably falling under the umbrella of regularization-based CL. While most of CL literature deals with image classification benchmarks, here we face the challenge of learning from unsupervised long visual experiences unrolling over time.

Unsupervised learning by motion cues. The idea of exploiting motion cues to support learning has been exploited for designing pretext tasks in the unsupervised learning community (Mahendran et al., 2019), aligning the similarity between pairs of feature vectors to the similarity between corresponding flow vectors or using segments (Pathak et al., 2017) from low-level motion-based grouping to train CNNs. Tiezzi, Marullo, Faggi et al. (2022) train models on video streams, leveraging a specific motion cue to ensure consistent representations of moving entities along a human-like attention trajectory (Tiezzi et al., 2020), that is given as an external signal. In contrast, CMOSFET employs a more sophisticated self-supervised objective that transcends the limitations of a simulated attention trajectory. Remarkably, our approach also learns to estimate the motion field (Marullo et al., 2022) without requiring any additional inputs. Our learning criterion is inspired by the concept of integrating motion prediction with feature extraction, while simultaneously acquiring these skills over time and at multiple levels of abstraction—a novel direction, to the best of our knowledge.

Learning in temporally-redundant streams. The very nature of videos, made of sequences of frames with slowly-varying, temporally-redundant signals, suggests intriguing research directions. Instead of

aiming at representing an individual frame, it is possible (Habibian, Yahia, Abati, Gavves, & Porikli, 2022) to estimate the difference between the representations of frames, so to ensure greater temporal consistency in the prediction output and faster processing in a knowledge distillation setting. In our work we do not estimate the feature deltas over time, but rather we assume that the extracted features are consistent when we take into account a motion-induced transformation between the frames. On a different note, unsupervised development of features from temporally coherent data has been investigated in Slow Feature Analysis (SFA) (Wiskott & Sejnowski, 2002), with more recent applications to high-level tasks, such as action recognition (Sun et al., 2014). The basic idea is to extract features that are slowly varying with respect to the quickly varying input signal, as it is the case for a natural video with moving entities. The concept of slowly varying features is similar to our temporally-consistent features, but in our work we do not impose any strong bias on their structure, they are defined at pixel-level and we explicitly bind their development to motion.

Contrastive self-supervised learning (SSL). In the context of unsupervised representation learning, unsupervised contrastive methods have recently closed the gap with supervised baselines in many vision tasks (He et al., 2020). Many existing approaches are based on contrasting the similarity of positive examples with the dissimilarity of negative examples (Chen, Kornblith et al., 2020). Recent works have replaced negative pairs with asymmetries in the way features are extracted (Chen & He, 2021; Grill et al., 2020), with still limited but improving theoretical understanding (Chen & He, 2021; Zhang et al., 2022). In order to increase the stability of the learning dynamics, many works involve an auxiliary model trained by moving average of the main network parameters (Grill et al., 2020; He et al., 2020), an intuition resumed by CMOSFET.

The vast majority of the methods can be considered global contrastive learning, since they are explicitly designed to obtain discriminative features for global tasks (e.g., image classification, or for self-supervised video representation learning (Li et al., 2021)). For tasks that require local details (e.g., semantic segmentation), there exist a few recent works on pixel-wise features (Xie et al., 2021; Xiong, Ren, Zeng, & Waabi, 2021), even though they train offline on large scale data or learn in a latent space at smaller resolution, thus not at a truly pixel-wise level. CMOSFET really embraces this challenge to learn pixel-wise representations from a single long video stream. Moreover, it works in a continual manner, which implies a completely different learning dynamics from traditional SSL methods. Those procedures typically exploit collections of small clips, offline-processed by stochastic optimization (Qian et al., 2021), and mainly for global video representation learning. Conversely, in our experiments we exploit the benchmark proposed by the main competitor (Tiezzi, Marullo, Faggi et al., 2022), which mimics lifelong video stream, showing that CMOSFET achieves the best results.

3. Model

Let us assume the availability of an endless collection of subsequent frames from a visual source, $\{I_t | t > 0\}$, originating a video stream \mathcal{V} at the resolution of $W \times H$, with \mathcal{X} the set of valid pixel coordinates for any frame. Each spatial location $x \in \mathcal{X}$ of I_t can be associated with a feature vector, carrying information about what is present in such a location and its neighborhood. We talk about *pixel-wise feature maps* when referring to the collection of such feature vectors. This paper focuses on the challenge of progressively developing a robust feature extractor \mathcal{F} leveraging the information in \mathcal{V} , working in a continual online manner, without storing buffers of past data, and without any supervision. The learnt features can then be used to preprocess visual stimuli in the context of different downstream tasks. It is important to remark that frames are continuously streamed at a constant frame rate, without any temporal limits (t could be potentially ∞) and they smoothly change over time, thus they are *not* independent. Under these

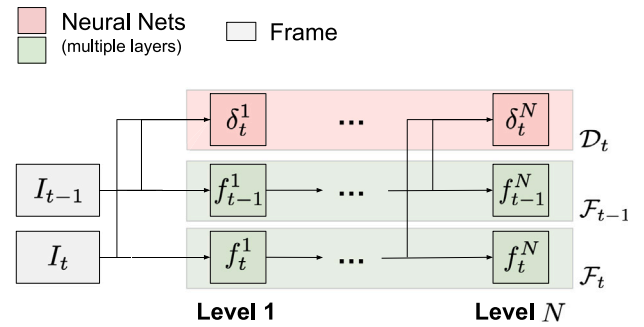


Fig. 1. The architecture of CMOSFET. Given a pair of consecutive frames I_{t-1} and I_t , pixel-wise features (f_{t-1}^ℓ, f_t^ℓ) and motion flow δ_t^ℓ are extracted at multiple levels, indexed by ℓ .

conditions, effective learning of useful features over time is not trivial and we will adopt ad-hoc components for the purpose. Following the deep and compositional nature of most modern neural architectures for computer vision, we assume that processing the stream \mathcal{V} with model \mathcal{F} results in extracting pixel-wise features at multiple ($N \geq 1$) levels of abstraction, indexed by $\ell \in [1, N]$. Another key notion for the purpose of this paper is the one of optical flow, that is widely known in the computer vision literature. Optical flow is the apparent motion associated to the pixels of the visual input, I_t , and it can be estimated with a variety of algorithms (Brox, Bruhn, Papenberger, & Weickert, 2004). Flow estimation can be efficiently performed by neural networks (Dosovitskiy et al., 2015), also when trained in an unsupervised manner (Teed & Deng, 2020), and a recent study (Marullo et al., 2022) showed that networks for flow extraction can be trained in a continual online fashion, without replay buffers or specific continual learning methodologies.

In the following we fully describe a model that, for each pixel of the input frame, computes both (i.) visual features and (ii.) motion estimation, at multiple levels of abstraction. Such information is strongly coupled, since the visual feature extractors and the motion estimators are learned in a joint manner. In particular, motion is the only driving signal for the development of the visual feature extractor, which conquers equivariance properties naturally induced by the motion fields. In the remainder of this Section, we describe how we achieve multi-order features and motion estimation in Section 3.1, the way in which such representations are linked and learned together in Section 3.2, the contrastive learning method we propose to avoid model collapse and trivial solutions in Section 3.3 and the way in which learning is enforced over time in Section 3.4.

3.1. Multi-order feature flows

We depict in Fig. 1 the architecture of the proposed CMOSFET. Given a pair of consecutive frames I_{t-1} and I_t , we denote with f_t^ℓ the feature extractor of the ℓ th level, working at time t , and consuming the output of $f_{t-1}^{\ell-1}$ as inputs as illustrated by green boxes in Fig. 1. The notation $f_t^\ell(x)$ indicates the feature vector of length d^ℓ from level ℓ at the 2D spatial coordinates x and $\mathcal{F}_t = \{f_t^\ell, \ell = 1, \dots, L\}$ with $\ell = 0$ being the degenerate case ($f_t^0 = I_t$). We now introduce the generic notion of *feature flow* as a generalization of optical flow when associated to arbitrary visual features f_t^ℓ , including I_t and more abstract features, motivated by the theoretical insights of Betti et al. (2021) and Betti, Gori et al. (2022). The notation δ_t^ℓ is used to indicate the function that extracts the ℓ th order feature flow, being $\delta_t^\ell(x)$ the flow vector at location x . Hereinafter the symbols f and δ will be overloaded depending on the specific context in which they are exploited, using them to denote functions or to represent their respective outputs, without any ambiguities.

It is convenient to represent motion in terms of displacement vectors between two consecutive frames, so that the feature flow δ_t establishes a point-wise correspondence between some features at f_{t-1} and f_t , and it allows the approximate reconstruction f_{t-1} given f_t (or vice-versa).¹ Formally, neglecting the superscript ℓ :

$$f_{t-1}(x) \approx W(f_t, \delta_t)(x) := f_t(x + \delta_t(x)) \quad (1)$$

being W the *warping* operator. As we depict in the red boxes of Fig. 1, the flow δ_t can be computed by a neural network that processes the representations of a pair of frames, at time $t-1$ and t , following Dosovitskiy et al. (2015) and Marullo et al. (2022). We cast this principle into the N -levels architecture described so far, so that δ_t^ℓ is the feature flow at level ℓ , and it exploits the frame representations produced at the layer below, $f_t^{\ell-1}$, where we define $f_t^0(x) := I_t(x)$ for any timestep, i.e., the original pixel representations. We use the notation \mathcal{D} to indicate the collection of all the neural networks δ^ℓ 's, dedicated to the computation of flows, to distinguish them from the feature extractors in \mathcal{F} .

The classic optical flow has a clear meaning related to cinematic apparent motion of objects with respect to the camera. However, it is largely known that this correspondence loses its strict physical interpretation when considering non-ideal conditions (e.g., strong occlusions, non-textured objects, changes of lightning, etc.) (Horn & Schunck, 1981). This consideration becomes very important when interpreting the feature flows δ_t^ℓ with $\ell \geq 1$. In deep architectures, higher-level features typically correspond to increasingly abstract representations that emerge from the learning process. As a result, it is challenging to explicitly define expectations for δ_t^ℓ . Additionally, higher-level features tend to encode information from a larger neighborhood around a given location in I_t (e.g., receptive fields in convolutional networks), causing local changes in I_t to potentially affect feature vectors associated with distant locations. Consequently, our multi-order flows are not accompanied by predetermined interpretation but rather serve as latent signals that facilitate the transfer of information among motion-connected locations, as we will delve into shortly.

3.2. Flow-conjugated representations

Several studies have investigated the idea of acquiring features that align with motion (Bregler, 1997; Li et al., 2021; Pathak et al., 2017; Xiong et al., 2021). Recently, Betti, Gori et al. (2022) introduced the notion of feature fields which are “conjugate fields with respect to motion” and specify the connections between a collection of arbitrary pixel-wise characteristics and the motion signal. We consider the bidirectional *conjugation* between features and their respective flows of Eq. (1) as an essential desirable property for learning purposes. On one hand, fulfillment of the approximation in Eq. (1) can be interpreted as the constraint that enforces flows to be consistent with respect to some “given features”, as in classic optical flow, where the features are brightness levels or colors. On the other hand, it stimulates “learnable features” on consecutive frames to be consistent with some estimated flow, that is the dual of plain optical flow, where features are given (visual input I_t). Formally, we define the following consistency penalty \mathcal{L}_c for the constraint of Eq. (1),

$$\mathcal{L}_c(\delta_t, f_{t-1}, f_t) = \frac{1}{wh} \sum_x \rho(f_{t-1}(x) - W(f_t, \delta_t)(x)), \quad (2)$$

being ρ a differentiable penalty function, that we selected to be the generalized Charbonnier photometric distance, $\rho(a) = (\|a\|^2 + \epsilon)^\zeta$ (Sun, Roth, & Black, 2010), with $\epsilon = 0.001$ and $\zeta = 0.5$, as in Marullo et al. (2022).

¹ We removed the layer index on purpose, since this description is generic. As it will become clear shortly, multiple feature levels can be associated to the same flow.

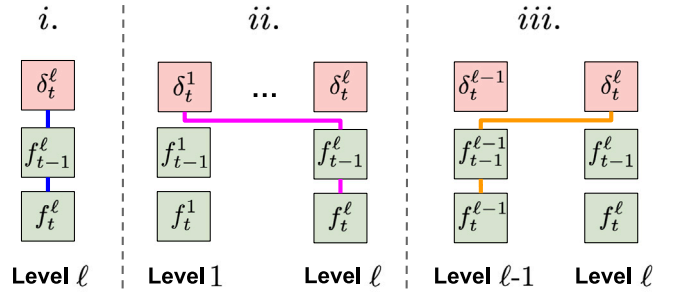


Fig. 2. Illustration of the three \mathcal{L}_c terms (i., ii., iii.) in Eq. (3). Each of the three sub-pictures include portions of the architecture of Fig. 1, while connections indicate the dependency introduced by the \mathcal{L}_c term. Case i. is about a single level, while ii. and iii. introduce cross-level dependencies.

Conjugation Loss. The consistency penalty of Eq. (2) represents a generic loss that could be straightforwardly exploited for learning purposes, using the features and motion yielded at each level. However, such implementation would not introduce any explicit cross-layer relation between features and flows (apart from the one arising from the cascaded feature computation). We propose to exploit three instances of the penalty in Eq. (2) to set up a loss function that shapes the way features and flows are developed across the model,

$$\begin{aligned} \mathcal{L}_{conj}^\ell = \mathcal{R}(\delta_t^\ell) & \\ & + \mathcal{L}_c(\delta_t^\ell, f_{t-1}^\ell, f_t^\ell) \quad \bullet (i.) \\ & + \mathcal{L}_c(\delta_t^\ell, f_{t-1}^1, f_t^1) \quad \bullet (ii.) \\ & + \mathcal{L}_c(\delta_t^\ell, f_{t-1}^{\ell-1}, f_t^{\ell-1}), \quad \bullet (iii.) \end{aligned} \quad (3)$$

where $\mathcal{R}(\delta_t^\ell)$ is the usual regularization term of optical flow, proportional to the sum of squares of the spatial gradient of the flow, $\nabla \delta_t^\ell$, and required for the well-position of the flow extraction problem (Horn & Schunck, 1981).² The role of the other three operands in the sum of Eq. (3) is qualitatively illustrated in Fig. 2, and described in the following: (i.) is the consistency between features and flow computed at layer ℓ , as expected; (ii.) is the consistency of the features of layer ℓ and motion from the first level, encouraging f^ℓ to be compatible with the displacement yielded by classic optical flow—i.e., learning promotes the preservation of a connection between higher-level features and the lowest-order motion; (iii.) is the consistency between motion estimated in level ℓ with the features of the previous level, to encourage higher-order motion flows to be compatible with less abstract features. We notice that the (regularized) term iii. is the usual loss for optical flow estimation applied at the different levels of the feature hierarchy. It is interesting to analyze the meaning of the conjugation loss in the case of δ_t^1 , the lowest-order flow. Intuitively, this flow is closely related to classic optical flow, since keeping only iii. in Eq. (3) makes it equivalent to the usual loss for optical flow estimation (recall that f^0 is I). Moreover, terms i. and ii. degenerate to the same term. Motivated by this considerations, we stop gradients to propagate to δ_t^1 through the i. and ii. losses, that allows our model to guarantee that δ_t^1 is coherent with the classic optical flow, that also drives the development of f_{t-1}^ℓ and f_t^ℓ for all the levels of the network.

3.3. Self-supervised contrastive learning

Despite these feature-flow constraints, the objective of Eq. (3) can be easily minimized by trivial solutions (e.g., spatially uniform features with arbitrary flow). To counteract this issue, we introduce a

² For the sake of readability, we omit the scaling factors in front of each term in the summation.

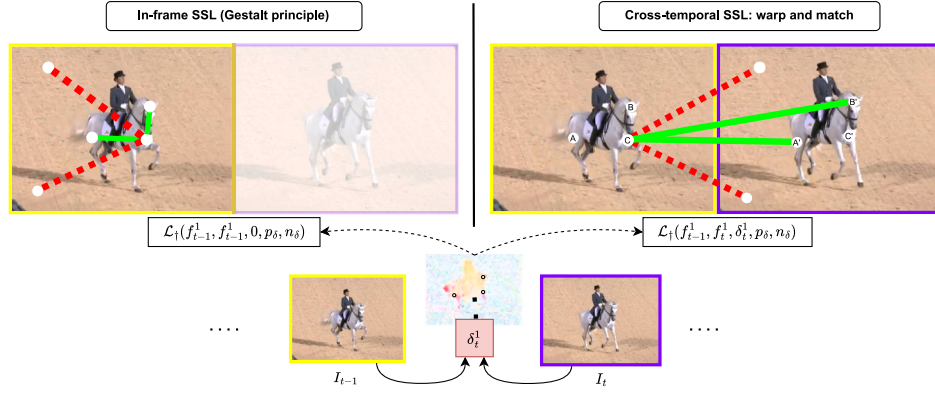


Fig. 3. Self-supervised loss, considering a pair of consecutive frames (I_{t-1}, I_t) , at the first level ($\ell = 1$) of the feature hierarchy (notice that it holds for any ℓ). The objective L_{self}^1 (Eq. (6)) is the sum of two contrastive penalties computed through L_1 . The first contribution (left) encourages the development of features by comparing features extracted on the same frame (I_{t-1}), while the second term (right) encourages alignment between features extracted in a pair of frames (I_{t-1}, I_t), thanks to flow matching (warping coordinates of items on one side of similarity/dissimilarity relationships, e.g. $A \rightarrow A'$, $B \rightarrow B'$, $C \rightarrow C'$). Green (red) links connect pixels whose features are enforced to be similar (different) according to our motion-based criterion.

flow-driven contrastive loss to enhance the emergence of meaningful features. Unlike popular contrastive approaches that are aimed at image-level global features (Chen, Kornblith et al., 2020; Liu et al., 2021), here we focus on the case of pixel-wise representations, thus we reconsider the role of positive and negative examples in contrastive learning. We propose to follow the Gestalt principle, which states that things that move together often carries similar semantic information, to establish relationships across consecutive frames when learning from a visual stream. In order to formally present our contrastive loss, we need to introduce three further ingredients, that are (i) a similarity function s , (ii) a way to determined positive and negative pairs, (iii) a sampling strategy to make learning affordable. We will avoid specifying the layer index ℓ unless explicitly needed.

(i) *Similarity Function.* The similarity function s is responsible of comparing features in a generic pair of pixel coordinates (x_i, x_k) . Such features are obtained by extractors that we will generically indicate with g and h for the first and second element of the pair, respectively. We introduce two different extractors since we both consider the case in which g and h operate on the same frame, and the case in which they operate on consecutive frames at $t-1$ and t . In the former case, no warping of the pixel coordinates will be involved, that is equivalent to consider a degenerate instance of the warping function exploiting δ that is 0 for all its inputs. In the latter case, x_k is warped by motion $\delta(x_k)$ before the feature extraction, i.e., we retrieve the location of the warped pixel x_k according to the estimated flow δ , before extracting features. Formally,

$$s(x_i, x_k, \delta, g, h) := \tau^{-1} \text{sim}(g(x_i), h(W(x_k, \delta(x_k))))$$

that is the τ -scaled cosine similarity with $\text{sim}(a, b) := a \cdot b / (\|a\| \|b\|)$ and $\|\cdot\|$ the Euclidean norm.

(ii) *Soft Assignment of Positive and Negative Pairs.* Gestalt assumption (Wertheimer, 1938) is exploited to identify positive and negative pairs, considering as positive those pairs of pixels that move coherently, and as negative those pairs that move in a different manner. Given a pixel at x , all the spatial coordinates of a given frame are evaluated in order to identify positive and negative examples. As a result, a pair that is composed of x and one of the *nearby* pixels with *similar motion* is a *positive pair*, while a *negative pair* is composed of x and one of the *distant* pixels with *different motion*. While this is a sensible learning signal (Liu et al., 2021; Tiezzi, Marullo, Faggi et al., 2022), we need to mitigate the effect of the natural violations of the assumption (for example, a non-rigid object will have different motion patterns in different parts of its surface, or there could be a static clone of a moving instance, etc.). We indicate with $p_\delta(x_i, x_j)$ the confidence score of (x_i, x_j) being a

positive pair, while $n_\delta(x_i, x_j)$ is the confidence score of (x_i, x_j) being a negative pair, each of them in $[0, 1]$. Scores n_δ and p_δ are computed as

$$n_\delta(x_i, x_k) = \left[\text{sim}(\delta(x_i), \delta(x_k)) \leq \tau_n \right] \frac{\|x_i - x_k\|}{\sqrt{H^2 + W^2}},$$

$$p_\delta(x_i, x_k) = \left[\text{sim}(\delta(x_i), \delta(x_k)) > \tau_p \right] \left(1 - \frac{\|x_i - x_k\|}{\sqrt{H^2 + W^2}} \right) \quad (4)$$

where $[\cdot]$ is 1 if the condition in brackets is true, otherwise it is 0. Thresholds $\tau_p, \tau_n \in [-1, 1]$ are selected to filter out pairs with uncertain similarities, with the condition $\tau_p \geq \tau_n$ which ensures that p and n are non-zero in a mutually-exclusive manner. The rightmost operands of the products in Eq. (4) are the distance between x_i and x_k scaled in $[0, 1]$ and 1 minus such a distance, respectively. Noticeably, motion direction is not significant for static points (motion vector is smaller than a fixed threshold τ_m). As such, we set a couple of exceptions to Eq. (4): a moving point and a static point are marked as dissimilar independently on their distance ($n_\delta(x_i, x_k) = 1, p_\delta(x_i, x_k) = 0$), while a pair of static points is neither similar nor dissimilar ($n_\delta(x_i, x_k) = p_\delta(x_i, x_k) = 0$). See also the qualitative example of Fig. 3, where positive pairs are connected by green lines, while red lines link negative pairs.

(iii) *Sampling.* The number of positive and negative pairs might easily become large, since it is quadratic in the number of pixels, making the evaluation of a contrastive criterion computationally costly. For this reason, we introduce a motion-driven sampling procedure to make learning affordable. This choice is loosely inspired by Tiezzi, Marullo, Faggi et al. (2022), where the authors proposed to randomly sub-sample the pixel coordinates in function of an externally computed focus of attention trajectory. Differently, we propose a selection criterion based on motion and representations, following the spirit on which our whole proposal is rooted. The outcome of this procedure is a set of η coordinates $\tilde{\mathcal{X}}_t \subset \mathcal{X}$, for each t , with a certain cardinality η , on which the proposed contrastive loss will be evaluated. Each pixel x at time t has probability $P_t(x)$ of being sampled, with $\sum_{x \in \mathcal{X}} P_t(x) = 1$. We compute $\{P_t(x), x \in \mathcal{X}\}$ ensuring that (i.) the probability of sampling a point in moving areas is the same as sampling in static areas, and that (ii.) the probability of sampling in areas where the j th feature has the strongest activation (absolute value) is the same for all j 's. The rationale behind this idea is that (i.) we want to ensure that in shots with limited motion the sampling process considers both static and moving areas in a balanced manner (the contrastive loss is built on motion information), and (ii.) we want to ensure that different visual patterns are covered in a balanced way by the sampling procedure. In fact, features are naturally associated to (usually) different patterns of the considered input. Formally, \mathcal{M} is the set of coordinates of moving pixels and $\bar{\mathcal{M}}$ are the coordinates of the static points. We compute




Fig. 4. Illustration of different sampling strategies ($\ell = 1$). Orange dots (bottom row) are points sampled from a visual stream (top row) in which a chair is moving in a static background that includes three smaller objects (pillow, laptop, teapot). First row: frame, estimated flow (different colors are about different directions), winning feature (different colors are about different winning features). Second row, left-to-right: plain uniform sampling, motion-biased sampling, and the proposed sampling driven by both motion and winning features. In the last case, the sampled coordinates cover both the moving chair and other details of the image in a balanced manner, while the first and second case give more emphasis to the uniform region (first, second) or the moving object (second).

S_j , $j = 1, \dots, d$, where set S_j collects the coordinates of the pixels for which the j th component of the learned representation is the largest one, i.e., $S_j = \{x_z : \arg \max |f_{j-1}(x_z)| = j\}$,³ where $\arg \max a$ is intended to return the index of vector a associated to the largest component and $|\cdot|$ is the element-wise absolute value. It is trivial to see that each pixel coordinate x ends up in belonging to one of the following $2d$ possible sets,

$$\begin{aligned} \mathcal{A}_j &= S_j \cap \mathcal{M}, \\ \mathcal{A}_{2j} &= S_j \cap \bar{\mathcal{M}}, \end{aligned} \quad \text{for } j = 1, \dots, d,$$

$$\text{where } \bigcup_{j=1}^d (\mathcal{A}_j \cup \mathcal{A}_{2j}) = \mathcal{X} \text{ and } \mathcal{A}_j \cap \mathcal{A}_{2j} = \emptyset \text{ for all } j.$$

If we let $P_i(x)$ return $1/(2d \cdot \text{card}(\mathcal{A}_z))$, being \mathcal{A}_z the set to which x belongs and $\text{card}(\mathcal{A}_z)$ its cardinality,⁴ we get a balanced distribution of the probabilities across the sets, $\sum_{x \in \mathcal{A}_1} P_i(x) = \dots = \sum_{x \in \mathcal{A}_{2d}} P_i(x) = \frac{1}{2d}$, that fulfills our initial requirements. Given $P_i(x)$, we sample $\eta > 1$ pixel coordinates for each t , collecting their coordinates into $\tilde{\mathcal{X}}_t$, that is smaller than \mathcal{X} . In Fig. 4 we report an example of our sampling strategy, comparing uniform sampling with motion-driven sampling and, finally, our motion-and-feature driven sampling. The picture shows how the proposed sampling procedure yields samples that are both related to moving objects (chair) and to portions of the frame with different visual patterns. Differently, a uniform sampling ends up in giving a lot of emphasis to the empty regions of the image. Fig. 3 highlights with white circles the elements of $\tilde{\mathcal{X}}_t$ (toy example).

Contrastive Loss. We are now ready to set up a contrastive loss in the usual log-exponential form (Le-Khac et al., 2020), giving more importance to positive pairs with large p_δ and to negative pairs with large n_δ ,

$$\begin{aligned} \mathcal{L}_\dagger(g, h, \delta) &= \\ - \sum_{u, v \in \tilde{\mathcal{X}}} \frac{p_\delta(u, v)}{Z} \log \frac{e^{s(u, v, \delta, g, h)}}{e^{s(u, v, \delta, g, h)} + \sum_{z \in \tilde{\mathcal{X}}} n_\delta(u, z) e^{s(u, z, \delta, g, h)}}, \end{aligned} \quad (5)$$

³ We use f_{i-1} to be consistent with the coordinate frame of δ_i , and $|a|$ is the element-wise absolute value of vector a .

⁴ The overall number of sets ($2d$) can be smaller in the case of one or more empty intersections.

being Z the normalization factor, $Z = \sum_{u, v \in \tilde{\mathcal{X}}} p_\delta(u, v)$. The outer summation is restricted to positive pairs, weighted by their degree of positiveness p_δ . The first element of each positive pair, u , acts as an anchor, so that the similarity of the positive pair is contrasted by all the negative pairs involving u (weighted by n_δ), as can be appreciated by the denominator of the log-argument. Notice that the sums of Eq. (5) are restricted to the sampled coordinates belonging to $\tilde{\mathcal{X}}_t$. Finally, our self-supervised objective L_{self}^ℓ is obtained by instantiating Eq. (5) twice, relating features from the same frame (at $t-1$)⁵ and across two consecutive frames ($t-1$ and t), respectively,

$$L_{self}^\ell = \mathcal{L}_\dagger(f_{t-1}^\ell, f_{t-1}^\ell, 0) + \mathcal{L}_\dagger(f_{t-1}^\ell, f_t^\ell, \delta_t^\ell). \quad (6)$$

leading to in-frame contrastive objective and a cross-temporal contrastive objective, respectively. Notice that, in the first term of the summation, the flow is set to 0, since no warping is performed, as already discussed.

3.4. Learning over time

Pairs of consecutive frames are processed in an online fashion, with a single-pass approach for each pair, and without any input from auxiliary memory buffers. Fig. 5 illustrates the proposed learning mechanism, CMOSFET, which is detailed in the following.

At each time step t , a new frame I_t is provided by \mathcal{V} and the previous one, I_{t-1} , is kept in cache. Learning consists in performing a single update of the model parameters in the whole network (F and D) given frames I_{t-1} and I_t , with the goal of minimizing the total loss L that considers both Eqs. (3) and (6),

$$L = \sum_{\ell} \left(L_{conj}^\ell + L_{self}^\ell \right). \quad (7)$$

Such a loss function L is naturally regularized over time, due to the L_{conj}^ℓ terms in the representation learning criteria, making it a natural instance of regularization-based approaches to continual learning. Moreover, the spatio-temporal bridge introduced on the sampled points by our contrastive loss (second term of the summation in Eq. (6)), introduces further regularity over time. Although it is possible to incorporate

⁵ This is due to the fact that motion predictors output displacement vectors in the reference frame at $t-1$.

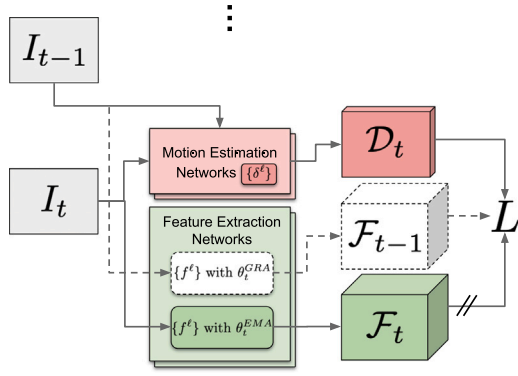


Fig. 5. Learning over time with CMOSFET. Exponential Moving Average (EMA) networks and gradient-updated networks (GRA) extract features from I_t and I_{t-1} , respectively. White-dotted parts of the model are in addition to the ones of Fig. 1, and the diagonal bars (//) indicate that no gradients are propagated on that path. The loss function L of Eq. (7) drives the learning process of the motion predictors and of the feature extractors, enforcing GRA to be coherent with EMA and instantiating the motion-induced contrastive criterion.

replays or other continuous learning techniques (Delange et al., 2021) alongside our proposed method (that goes beyond the scope of this paper), this study focuses exclusively on the more demanding yet more authentic scenario of plain continual online learning (Tiezzi, Marullo, Faggi et al., 2022). It is natural to wonder whether such regularity over time is enough to guarantee a good compromise between plasticity and stability of the networks, hence mitigating catastrophic forgetting, in both the motion estimators D and the feature extractors F (red and green boxes in Fig. 5, respectively).

A recent experience on motion estimation in continual learning settings (Marullo et al., 2022) has shown that the (regularized) criterion of Eq. (2) is enough to learn to estimate pixel displacements δ over time by online gradient descent. Moreover, forgetting issues are not particularly evident due to the locality of the motion prediction process, resulting in negligible interference even in case of long streams with significant variability. However, our initial experimental findings indicated that learning over time poses significant challenges for the feature extractors F , as we observed notable forgetting when approaching the problem using the naive approach. Then, we borrow an intuition from models whose weights are updated by a momentum-based moving average (Arani, Sarfraz, & Zonooz, 2022; Cai et al., 2021; He et al., 2020; Tarvainen & Valpola, 2017). Despite being conceptualized in different configurations, these approaches share the idea of using an Exponential Moving Average (EMA) scheme to update the parameters of a (teacher) network, while another network (student), which is continuously updated by gradient descent, is enforced to be consistent with it, preventing rapid changes in the parameter space. The EMA-updated network acts as a slowly progressing encoder, with parameters that evolve smoothly (He et al., 2020), and it can be employed both to stabilize learning of the student network and to better preserve information from the past.

Since our model is designed to extract and relate representations from a pair of frames (I_{t-1} , I_t), we propose to extract features from frame I_{t-1} using a feature extractor whose weights at time t , referred to as θ_t^{GRA} (GRADient-updated), are updated by an online gradient step, while the current frame is processed by an EMA-updated network (slow learner) with the same architecture and weights θ_t^{EMA} ,

$$\begin{aligned} f_{t-1}^\ell &= f^\ell(f_{t-1}^{\ell-1}, \theta_t^{GRA,\ell}) \\ f_t^\ell &= f^\ell(f_t^{\ell-1}, \theta_t^{EMA,\ell}), \end{aligned}$$

where, for the sake of simplicity, we overloaded the notation used so far making explicit the weight values-related argument. Fig. 5 summarizes the structure of whole model, emphasizing the two types of networks GRA and EMA.

Such two networks are naturally bridged by the L_{self} (Eq. (6)) component of the loss function L , that enforces coherence between them in the motion-driven manner proposed in this paper. While the GRA network is updated by the gradient of L with respect to its weights and using learning rate $\alpha_f > 0$, the other network is updated by EMA with coefficient $\xi \in [0, 1)$,

$$\begin{aligned} \theta_{t+1}^{GRA} &= \theta_t^{GRA} - \alpha_f \nabla_{\theta} L(F_{t-1}, \hat{F}_t, D_t) \\ \theta_{t+1}^{EMA} &= \xi \theta_t^{EMA} + (1 - \xi) \theta_{t+1}^{GRA}, \end{aligned}$$

where $\nabla_{\theta} L$ is the gradient with respect to the weights in F and \hat{F}_t indicates that F_t is treated here as a constant value (it comes from the EMA instance). The learning procedure is summarized in Algorithm 1.

Algorithm 1 The learning procedure. For each frame from stream \mathcal{V} , features are extracted, motion is estimated, and the proposed self-supervised loss L drives the learning processes. The GRA and EMA networks are initialized to the same weight values.

Require: Stream \mathcal{V} of length T , potentially infinite; neural nets D and F with initial parameters set to γ_1 and θ_1 . Learning rates α_m , $\alpha_f > 0$ and $\xi \in [0, 1)$.

$$t \leftarrow 1, \quad \theta_1^{GRA} \leftarrow \theta_1, \quad \theta_1^{EMA} \leftarrow \theta_1$$

$$f_0^0 = I_1, \quad f_1^0 = I_1$$

while $t \leq T$ **do**

$$I_t \leftarrow \text{next_frame}(\mathcal{V})$$

for $\ell = 1, \dots, \ell = N$ **do**

$$\delta_t^\ell = \delta^\ell(f_{t-1}^{\ell-1}, f_{t-1}^{\ell-1}, \gamma_t^\ell)$$

$$f_{t-1}^\ell = f^\ell(f_{t-1}^{\ell-1}, \theta_t^{GRA,\ell})$$

$$f_t^\ell = f^\ell(f_t^{\ell-1}, \theta_t^{EMA,\ell})$$

end for

$$D_t \leftarrow \{\delta_t^\ell | \ell = 1, 2, \dots, N\}$$

$$F_{t-1} \leftarrow \{f_{t-1}^\ell | \ell = 1, 2, \dots, N\}$$

$$F_t \leftarrow \{f_t^\ell | \ell = 1, 2, \dots, N\}$$

$$\gamma_{t+1} \leftarrow \gamma_t - \alpha_m \nabla_{\gamma} L(F_{t-1}, F_t, D_t)$$

$$\theta_{t+1}^{GRA} \leftarrow \theta_t^{GRA} - \alpha_f \nabla_{\theta} L(F_{t-1}, F_t, D_t)$$

$$\theta_{t+1}^{EMA} \leftarrow \xi \theta_t^{EMA} + (1 - \xi) \theta_{t+1}^{GRA}$$

$$t \leftarrow t + 1$$

end while

4. Experiments

The proposed framework, CMOSFET, has been developed using the PyTorch library, running experiments on two Linux machines equipped with an NVIDIA Tesla V100 GPU, 32 GB of memory (see Appendix for further details). The primary objective of this work is to explore the potential of CMOSFET in generating pixel-wise representations, learning in a continual, online, self-supervised manner, concurrently developing multiple flows estimators. Our experimental evaluation follows a recently proposed benchmark that was proposed in Tiezzi, Marullo, Faggi et al. (2022) and also inherited by further works (Tiezzi, Marullo, Betti et al., 2024), specifically designed for evaluating the quality of pixel-wise features generated in a continual learning paradigm, that we summarize in the following, after having introduced the datasets we selected.

Datasets. Our proposed method, CMOSFET is evaluated on two distinct collections of video streams. The first collection, based on the work of Tiezzi, Marullo, Faggi et al. (2022), comprises three 256×256 streams generated from the 3D Virtual Environment SAILenv (Meloni et al., 2021; Meloni, Pasqualini, Tiezzi, Gori, & Melacci, 2020). Video streams are about some target objects moving around and performing complex transformations consisting of rotations, scaling, and pose variations with respect to a fixed camera. The first stream, EMPTYSPACE, depicts a room with uniform background and four moving objects, including a chair, laptop, pillow, and ewer. It includes both grayscale

(-BW suffix) and RGB (-RGB) data. The second stream, `SOLID`, comprises three white 3D solids, including a cube, a cylinder, and a sphere, moving in a grayscale environment. This stream requires the developed features to discern context/shapes rather than brightness or color. The third stream, `LIVINGROOM` (-BW, -RGB), depicts the same four objects as in `EMPTYSPACE`, moving around in a more complex scene with a heterogeneous background composed of non-target objects such as a couch, tables, staircase, door, and floor, as well as static copies of the `EMPTYSPACE` objects. We inherit from [Tiezzi, Marullo, Betti et al. \(2024\)](#) the second collection of video streams, which include two real-world videos, `RAT` (256×128) and `HORSE` (256×192), proposed by [Liang, Li, Jafari, and Chen \(2020\)](#), that recorded the behavior of rat and a horse (actually horse+jockey), respectively, representing the target objects of the streams. These videos⁶ are used to demonstrate CMOSFET’s ability to generalize to real-world scenarios, with non-fixed camera.

Experimental Setup. In order to evaluate the performance of CMOSFET, we conducted a comparative study with the main competitor, the model proposed in [Tiezzi, Marullo, Faggi et al. \(2022\)](#), on a pixel-wise classification task. Our experiments strictly adhere to the benchmark framework of [Tiezzi, Marullo, Faggi et al. \(2022\)](#), which consists of unsupervised learning for 30 laps⁷ per object, with the last 5 laps involving the storage of 3 supervised representations (pixel-wise templates) per object, at intervals of 100 frames, in a small memory referred to as C . These externally provided cues are only utilized for an evaluation (classification) procedure, which has no bearing on the representation learning process. They simulate the outcome of a tiny set of interactions with a supervisor, as is typical of realistic learning settings. The evaluation stage consists in processing the stream for one more lap, not only extracting features but also providing them to an open-set classifier $c(\{f_i^\ell(x)\}_{\ell=1, \dots, N}, C)$ that predicts the class-membership scores of pixel coordinates x given the representations produced by feature extractors in \mathcal{F} and the templates stored in the memory C . Scores are about a certain number of classes that depends on the considered stream. The quality of the predictions of c is measured by the F1-score. The classifier c is implemented as a distance-based model (cosine similarity in our case).⁸ In such a way, the model is capable not only to predict the target classes, but also to abstain⁹ from prediction when the minimum distance from all the templates populating C is greater than a certain threshold ξ . In the case of CMOSFET, we apply a simple learning schedule to warm-up the motion estimators, avoiding premature training of the neural feature extractors. During the first T_{sched} laps we only learn the flow estimator at level 1. Then, if $N > 1$, we start learning the feature extractor on the same level, and we wait further T_{sched} laps before applying the same schedule on the following level. The procedure is repeated, progressively activating learning on the whole network (see [Appendix A.2](#) for details).

Following the evaluation protocol of [Tiezzi, Marullo, Betti et al. \(2024\)](#) and [Tiezzi, Marullo, Faggi et al. \(2022\)](#), the performance of our approach is evaluated by computing the F1 score (averaged over the available object categories and the background class) over all pixels of

the frames, in a last additional lap per object. We adopted the exact same methodology to evaluate the performance of our approach on the considered real-world videos. In this case, ground truth for evaluation comes from manual labeling provided by [Liang et al. \(2020\)](#) on selected frames. Still following the procedure defined in the evaluation protocol of [Tiezzi, Marullo, Betti et al. \(2024\)](#) and [Tiezzi, Marullo, Faggi et al. \(2022\)](#), the optimal values for the hyperparameters are chosen by maximizing the F1 along the trajectory (1 pixel per frame) of a given human-like attention model ([Zanca, Gori, Melacci, & Rufa, 2020](#)). The grids of parameter values that we considered and the optimal values are reported in [Appendix A.2](#), together with further details.

Neural Architectures. We now present the neural architectures used to implement the feature extractors \mathcal{F} and flow estimators \mathcal{D} . Each f^ℓ and each δ^ℓ consists of a RESUNET architecture ([Tiezzi, Marullo, Faggi et al., 2022](#)), which is a variant of the U-Net network ([Ronneberger, Fischer, & Brox, 2015](#)). We reduced the number of convolutional filters in each layer by a factor of 4, while the number of per-pixel outputs is 32 for f^ℓ and 2 for δ^ℓ (horizontal and vertical flow components). We considered two implementations of this model: CMOSFET-Single, with a single level of feature and flow extractors (i.e., $N = 1$), and CMOSFET, with two levels, developing lower and higher-order motions (i.e., $N = 2$). We compared with the competitors presented in [Tiezzi, Marullo, Faggi et al. \(2022\)](#): ResUnet and FCN-ND (which has 6 convolutional layers and no downsampling). For more details refer to [Appendix A.1](#). As a rough baseline for comparison, we also evaluated the performance of several models that were pre-trained and specialized in semantic segmentation tasks and self-supervised learning, of course without attempting to outperform their results. Specifically, we evaluated the ResNet101-based DEEPLABV3 ([Chen, Papandreou, Kokkinos, Murphy, & Yuille, 2017](#)) and the Dense Prediction Transformer - DPT ([Ranftl, Bochkovskiy, & Koltun, 2021](#)), both of which have backbones trained on ImageNet-1M (-B suffix) and classification heads (-C suffix) trained on COCO ([Lin et al., 2014](#)) and ADE20k ([Zhou et al., 2019](#)). We also considered the pre-trained features from the backbones of recent self-supervised approaches, including MoCo v1, v2, v3 ([Chen, Fan, Girshick and He, 2020](#); [Chen*, Xie*, & He, 2021](#); [He et al., 2020](#)) and PIXPRO ([Xie et al., 2021](#)). We upsampled the features picked at one of the last three residual stages (testing all of them) and reported the best-performing configuration. Additionally, we used the dummy baseline RAW IMAGE, which uses the original pixel representations (i.e., pixel color/brightness) as features. Concerning PIXPRO, we also tried learning the weights, applying the scheme proposed by the authors, in our lifelong online context. Since it was not designed to learn over time (it is an offline learning method), results are not competitive, probably due to the fact that such framework lacks an explicit use of a spatio-temporal contrastive loss.

Quantitative Results. [Table 1](#) reports the evaluation results of our proposed method, considering the F1 scores over the entire frame. Specifically, we focus on the related *continual* self-supervised competitors (bottom part of the table), observing that CMOSFET outperforms them in all video streams. The smaller variant CMOSFET-Single is already slightly superior to the competitors, except for `EMPTYSPACE-BW`, where it is on par with the best competitor. Notably, the largest gap between CMOSFET-Single and CMOSFET is observed in real-world videos (`RAT`, `HORSE`) that confirms the importance of higher-order motion. We remark that in our proposed approach, the flow estimation is learned online from scratch, in contrast to the competitors, that rely on a pre-computed motion signal that is flawlessly produced by the rendering engine. Moreover, CMOSFET methods do not rely on explicit segmentation procedures for learning, yet they effectively exploits the relationships between flows and learned features. Comparing `EMPTYSPACE-BW` and `EMPTYSPACE-RGB`, we notice a significant gap in F1, suggesting that even though the RAW scores of grayscale and color representations are not so different, both CMOSFET and the competitors seem to efficiently exploit the intrinsically richer information encoded in RGB. Additionally, we observe that CMOSFET encodes pixels in a

⁶ <https://www.kaggle.com/datasets/gvclsu/long-videos>.

⁷ The agent observes the scene from a fixed position, while individual objects of interest move one at a time along smooth trajectories. These objects rotate and change their distance from the camera, either moving closer or farther away. A complete route made by each object, returning to its starting point, is referred to as a *lap*. Please refer to [Tiezzi, Marullo, Faggi et al. \(2022\)](#) for further details on the experimental framework.

⁸ Features f_i^ℓ are independently normalized for each ℓ , in order to account for large differences in magnitude between the extractors.

⁹ Open-set classifiers ([Scheirer, Rezende Rocha, Sapkota, & Boulton, 2012](#)), in addition to being capable to distinguish between examples belonging to different training classes, can also detect whether data do not belong to any known class, a realistic condition for autonomous visual agents. CMOSFET is also class-incremental ([Geng, Huang, & Chen, 2020](#)), due to the progressive inclusion of new classes after human intervention.

Table 1

F1 scores over 10 runs (mean \pm std), in seven different video streams (EMPTYSPACE, SOLID, LIVINGROOM, RAT and HORSE with some -BW variants). The bottom part of the table is about the main competitors of the proposed model, including the raw-image (degenerate) baseline. The top part of the table collects reference results obtained with large-offline-pre-trained models, publicly available (with no attempts to overcome them).

	# Params	EMPTYSPACE		SOLID	LIVINGROOM		RAT	HORSE	
		BW	RGB	BW	BW	RGB	RGB	RGB	
OFFLINE PRE-TRAIN.	DPT-C (Ranftl et al., 2021)	121M	0.66	0.67	0.64	0.35	0.39	0.59	0.83
	DPT-B (Ranftl et al., 2021)	120M	0.71	0.69	0.68	0.39	0.39	0.58	0.87
	DEEPLAB-C (Chen et al., 2017)	58.6M	0.49	0.61	0.57	0.31	0.34	0.56	0.86
	DEEPLAB-B (Chen et al., 2017)	42.5M	0.70	0.65	0.66	0.34	0.44	0.57	0.81
	MoCo v1 (He et al., 2020)	8.5M	0.73	0.73	0.74	0.33	0.35	0.70	0.66
	MoCo v2 (Chen Fan et al., 2020)	8.5M	0.75	0.76	0.74	0.41	0.43	0.59	0.79
	MoCo v3 (Chen* et al., 2021)	8.5M	0.76	0.76	0.76	0.41	0.44	0.58	0.64
	PIXPRO (Xie et al., 2021)	8.5M	0.31	0.46	0.41	0.30	0.22	0.59	0.70
CONTINUAL.	RAW IMAGE	–	0.50	0.45	0.18	0.10	0.23	0.67	0.66
	HOURLASS (Tiezzi, Marullo, Faggi et al., 2022)	17.8M	0.55 \pm 0.03	0.71 \pm 0.03	0.50 \pm 0.01	0.31 \pm 0.04	0.25 \pm 0.07	0.59 \pm 0.04	0.71 \pm 0.07
	FCN-ND (Tiezzi, Marullo, Faggi et al., 2022)	0.1M	0.60 \pm 0.05	0.51 \pm 0.07	0.48 \pm 0.03	0.24 \pm 0.01	0.28 \pm 0.03	0.42 \pm 0.05	0.50 \pm 0.08
	PIXPRO (Xie et al., 2021)	8.5M	0.45 \pm 0.15	0.47 \pm 0.06	0.19 \pm 0.08	0.12 \pm 0.09	0.25 \pm 0.14	0.37 \pm 0.11	0.42 \pm 0.07
	CMOSFET-Single	1.1M	0.60 \pm 0.06	0.78 \pm 0.06	0.62 \pm 0.02	0.33 \pm 0.03	0.34 \pm 0.05	0.61 \pm 0.07	0.75 \pm 0.06
	CMOSFET	2.3M	0.64 \pm 0.06	0.82 \pm 0.02	0.64 \pm 0.05	0.38 \pm 0.04	0.36 \pm 0.05	0.74 \pm 0.09	0.84 \pm 0.01

reduced number of output features compared to competitors (32 · 2 vs. up-to –128), and has significantly fewer learnable parameters than the best competitor (on average), with 2.3M parameters compared to 17.8M (third column of Table 1). This confirms the capability of our method to develop informed but more compact representations. Surprisingly, our approach achieves interesting performance even when compared to offline-pre-trained large architectures with several millions of parameters (upper part of the table). It is important to note that our goal is not to overcome large-scale massively offline-trained models, since we learn from scratch in an online manner, specializing in the target environment. Nonetheless, CMOSFET beats these competitors in EMPTYSPACE-RGB and RAT-RGB. We highlight the performance on the other real-world video stream, HORSE-RGB, where CMOSFET overcomes most of the offline-trained models (all of the self-supervised ones) and it is only slightly behind the best one. These results affirm the adaptability of our approach to the properties of the environment at hand, with a fraction of the learnable parameters and processed data.

Qualitative Results. Fig. 6 presents a qualitative analysis of the outputs of our models on three sample frames, with similar trends observed across all video streams. Specifically, the second and third columns depict the extracted flows from CMOSFET’s first and second levels, respectively. The features produced by CMOSFET allow the classification procedure (recall that such procedure do not affect the feature learning process) to completely discriminate all the different target objects (fourth column, predictions). While the borders of the CMOSFET-based predictions appear slightly thicker compared to the ground-truth object borders, this is primarily due to the quality of the estimated motion field, which is not always perfect (recall that it is learned from scratch jointly with the features). In fact, contrastive learning is affected by noisy flow, which can easily result in false positive pairs, with one pixel belonging to an object and the other just outside of it. Noticeably, the estimated optical flow (second column) is very crisp in the case of EMPTYSPACE where the camera is fixed, while it captures the effects of the moving-camera in real-world streams (second and third row). Also, optical flow estimation is known to be inherently challenging in nearly texture-free surfaces. Interestingly, comparing the two flow estimators of distinct levels (second and third column), we notice how they appear to focus on different parts of the frames, with flows that are somewhat related but clearly not overlapping. This confirms that the network develop a latent higher-order motion, challenging to interpret, but well suited to the learning process. It is worth noting that we did not use any specific tricks (Zhai, Xiang, Lv, & Kong, 2021) to improve the flow extraction, such as spatial pyramids or occlusion handling, for the sake of simplicity, consistently with Marullo et al. (2022).

Ablation studies. In order to better understand the impact of key design choices, we perform ablation experiments concerning some of the main components of CMOSFET, reporting our findings in Fig. 7 that we describe from left to right, top to bottom. The motion-feature conjugation term L_{conj} from the loss function (Eq. (3)) provides improvements in all the streams (1st barplot), when compared to the case in which we simply enforce consistency with optical flow (i.e., Eq. (3) is composed of *ii.* only), more evidently in real-world streams. It also improves stability among the different runs, resulting in reduced standard deviation—apart for RAT, where the lack of full conjugation leads to collapse to a very suboptimal solution with minimal variance. The adoption of a specific sampling strategy (2nd barplot) for the contrastive term, guided by both Motion and Features (M.F. in Fig. 7), also had a considerable influence on the results, sometimes dramatic as in the case of RAT. In general, we notice that different sampling strategy might be optimal in different streams, in relation with the kind of motion that is present in the respective environments. Motion-only-guided sampling (Mot.) can yield slight improvements w.r.t. vanilla uniform sampling strategy (Plain), due to the fact that it helps focusing on the rather small moving areas. Learning without the stabilizing effect of the EMA network (3rd barplot) generally leads to lower performance, in agreement with literature in both contrastive learning and continual learning. The learning process is however still effective (apart from the challenging RAT), thanks to regularization effects introduced by the spatio-temporal loss terms, although less stable than when using the EMA scheme. We also separately evaluated the impact of the number of sampled locations η (4th barplot). It turns out that models developed in real-world streams consistently benefit from a larger number of locations ($\eta = 200$). Conversely, this trend is not evident in synthetic streams, that have either large static or uniform areas, where more dense sampling does not yield more information.

Limitations. Recent works highlighted the inherent unstable nature of self-supervised contrastive criteria (Zhang et al., 0000). CMOSFET is a pixel-wise motion-based contrastive criterion, proposed in the context of a completely online learning procedure, and as such there are some stability issues, confirmed by the non-negligible standard deviation of our results over multiple different initializations (Table 1). Additionally, contrastive methods can suffer from dimensional collapse (Jing, Vincent, LeCun, & Tian, 0000), where learned embeddings fall into a lower-dimensional subspace, with limited discriminative power. While our motion-based constraints partially mitigate this issue, the risk of collapse remains to be considered. Furthermore, effectiveness of contrastive learning relies heavily on well-constructed negative pairs. In dynamic environments with subtle or overlapping motion patterns, it can be challenging to identify hard negatives that do not overlap with positive samples, which may impact the robustness of the learned




Fig. 6. Three frames sampled from EMPTYSPACE, RAT and HORSE streams (first column). CMOSFET flows (first-order and higher-order in second and third columns, respectively) are plotted according to the optical flow conventional color mapping (Baker et al., 2007) (see the color wheel on the right). Predictions in the fourth column show a pretty satisfying adherence to the borders of the objects (target objects are highlighted with colors; the background class is present in *all* the streams, reported with a light gray overlay).




Fig. 7. Ablation of the model components on four visually-rich streams (L stands for LIVINGROOM, HOR for HORSE). From left to right: the impact of the features-motion conjugation term L_{conj} as presented in Eq. (3) (Full) compared to the case in which features are only developed by enforcing them to be coherent with optical flow (i.e., when Eq. (3) is composed of *ii.* only-Flow only); the role of different sampling strategies (Plain, Motion-Mot., Motion and Features-M.F.); EMA network (without or with); number of sampled points η .

representations. Another intrinsic source of instability stands on the sensibility to the effects of the first learning steps, where the flow estimators are obviously far from being able to make reasonable predictors. In fact, the proposed contrastive loss (Eq. (6)) depends on how motion gets predicted, thus early-stage prediction might drive the feature extractor(s) toward configurations that we sometimes found to have negative effects on the following development of the model. CMOSFET (as well as its main competitor) is designed to learn in environments where background areas are characterized by motion patterns that are different from the ones of objects to which semantics are expected to be attached (or that are almost static). Our results show that small (or relatively slow) camera motion, as the one of the selected real-world videos, can be easily handled or filtered out using a larger τ_m or setting it to the average length of motion vectors in the current frame, as we did in our implementation. However, dealing with strongly moving cameras might lead to development of features that encode information which is less object-specific. On the other hand, being this model completely based upon motion information, the lack of movement would be detrimental. In the trivial scenario where both the foreground and background objects move uniformly throughout the stream or have no movement at all, CMOSFET would not yield disentangled features. Notably, this phenomenon is also observed in humans—for instance, infants tend to oversegment static objects (or those moving identically) but can effectively group objects characterized by different motions (see Ostrovsky et al. (2009) and Spelke (1990)). Finally, the motion-driven sampling reduces computational costs but may introduce biases tied to observed motion patterns, potentially limiting generalizability. Using an higher number of pixels could help devise more robust features but at a higher cost.

5. Conclusions

In this work we focused on learning to predict pixel-wise information in a fully unsupervised continual manner. We considered multi-level features and multi-level (order) motion flows as “conjugated” one to each other, mutually constraining their development. We proposed CMOSFET, that, to our best knowledge, is the first implementation of these principles. CMOSFET is based on neural models trained to extract features and flows from scratch, driven by multi-level spatio-temporal feature-motion constraints and on a self-supervised (contrastive) loss that is biased by the way flows are predicted. We performed experiments on video streams from a recent benchmark and real-world videos, showing how CMOSFET benefits from learning multi-order flows, and how it significantly overcomes its main competitor. CMOSFET also behaves similarly to larger networks that were offline-pretrained on large collections of supervised and unsupervised images. The novelty of the ideas in this paper are paired with limitations which we explicitly discussed, and that will be the subject of our future studies. Finally, another direction for future work is the use of CMOSFET jointly with more advanced continual learning strategies to avoid forgetting issues in significantly longer streams or with several object categories.

CRedit authorship contribution statement

Simone Marullo: Writing – review & editing, Software, Methodology, Investigation, Data curation, Conceptualization. **Matteo Tiezzi:** Writing – review & editing, Software, Methodology, Investigation, Data curation, Conceptualization. **Marco Gori:** Writing – review & editing,

Formal analysis, Conceptualization. **Stefano Melacci**: Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partly supported by the University of Siena (Piano per lo Sviluppo della Ricerca - PSR 2024, F-NEW FRONTIERS 2024), under the project “Time-driveN StatEful Lifelong Learning” (TINSELL) and also by the project “CONSTR: a Collectionless-based Neuro-Symbolic Theory for learning and Reasoning”, PARTENARIATO ESTESO “Future Artificial Intelligence Research - FAIR”, SPOKE 1 “Human-Centered AI” Università di Pisa, “NextGenerationEU”, CUP I53C22001380006.

Appendix

This appendix provides further details about the proposed model. This includes in-depth descriptions of the neural architectures and the selected competitors, the grids of values of the hyper-parameters that we validated in our experimental activity, together with the best selected values. We also provide additional results and qualitative investigations, as well as further information on how to reproduce the results of this paper (including implementation-oriented descriptions). Notice that the bibliographic references contained in this appendix are about the bibliography of the main paper.

A.1. Neural architectures

We provide additional details on the selected neural architectures and on the competitors.

Feature Extractors. The neural branch of CMOSFET-Single yielding features, i.e., function f , is a ResUnet architecture, inspired to the one presented in Tiezzi, Marullo, Faggi et al. (2022) but having a smaller amount of parameters. It is a UNet-like (Ronneberger et al., 2015) architecture,¹⁰ based on a ResNet18 backbone. Among the pretrained competitors, DPT-Hybrid (Ranftl et al., 2021) is a Transformer-based ViT model, trained on the ADE20K dataset (Zhou et al., 2019). All the other considered models adopt ResNet backbones, in particular the ResNet-50 architecture, except for DeepLab v3 (ResNet-101). DeepLab models apply Atrous Spatial Pyramid Pooling (ASPP) on top of the backbone features and they are trained on COCO dataset (Lin et al., 2014).¹¹ MoCo models¹² are trained with a contrastive unsupervised objective on ImageNet-1M. MoCo v2 differs from the original version because of an improved training procedure (better augmentation, cosine learning rate schedule, extra MLP in training phase). MoCo v3¹³ introduces some tricks to improve the learning stability both for ViT and Resnets backbones. We focus our analysis on the latter backbone. In the case of PixPro¹⁴ models, we use backbone weights obtained from pretraining on ImageNet-1M with their pixel-level contrastive task.

Motion Estimation. CMOSFET neural branch for motion estimation, i.e., function δ , inherits its structure and composition from the

feature extraction one (ResUnet), except that we cut the skip connections connecting the input to the output layer, since we empirically noticed in preliminary experiments that we could obtain a more sensible estimation. Such a network outputs the vertical and horizontal displacement maps for all the frame pixels.

A.2. Additional implementation details and hyper-parameters

In order to complement the already described experimental setup of Section 4, here we introduce some further details to make our experience reproducible. For completeness we report that, in our implementation, the maximum distance between pairs of pixels in Eq. (4) (i.e., $\sqrt{H^2 + W^2}$ in the normalization factors) is replaced with the maximum distance between pairs of sampled coordinates.

Data augmentation. For each processed frame pair, we create a mini-batch composed of the two original frames (I_{t-1}, I_t), as well as multiple augmented views/transformations, in order to improve the feature robustness, as common in deep net training procedures. Augmentations are of 3 different types: (A) large random crops (with crop ratio greater than 0.9) followed by upscaling to $W \times H$; (B) random horizontal/vertical flips; (C) random color distortion and Gaussian blur. For the augmentations of type (A) and (C), only one of the frames composing the pair (I_{t-1}, I_t) is transformed, randomly selected at each t , whilst the other is kept as it is.

Parameters and validation. In order to make the comparisons fair, we followed the exact same hyper-parameter validation procedure of Tiezzi, Marullo, Faggi et al. (2022). In particular, for each video stream we selected the values of the hyper-parameters which maximize the F1 score measured during the 30th lap, only considering one pixel per frame, along the trajectory of a human-like focus of attention mechanism attending the scene. We sampled the best hyper-parameters from the following grids: $\lambda_m \in \{1\}$, $\beta_m \in \{10^{-4}, 10^{-3}, 10^{-2}\}$, $\beta_f \in \{1\}$, $\tau_p \in \{0.7, 0.8, 0.9\}$, $\tau_n \in \{-0.5, -0.3, 0.3, 0.5\}$, $\ell \in \{50, 100, 200\}$, $\tau \in \{0.1, 0.5\}$, $\tau_m \in \{0.5, 1.0, 1.5, 2.0\}$, $\alpha_m \in \{10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 10^{-2}\}$, $\alpha_f \in \{10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 10^{-2}\}$, $\xi \in \{0.5, 0.9, 0.99\}$.

An additional speed-up of the training process can be achieved by further limiting the summations of Eq. (5) (both of them) by keeping only a fraction $\aleph \in (0, 1]$ of the positive and negative pairs (balanced). In particular, we experimented the case in which we kept those positive (resp. negative) pairs for which representations are the least similar (resp. most similar), i.e., focusing on the ones that mostly contribute to larger values of L_{self} . In the hyperparameter search, we consider $\aleph \in \{0.5, 0.7, 1.0\}$. The motion estimator δ is optimized using Adam (Kingma & Ba, 2014), while we empirically found better results by optimizing the features extractor f with SGD (with the exception of the EMPTYSPACE stream, where Adam usually yields better performances). In addition to the typical smoothness regularizer introduced in Eq. (3), we also exploited a further regularization modulated by a customizable parameter λ_r , in order to keep the estimated motion magnitude small in the case of uniform backgrounds,

$$\mathcal{R}(\delta_t^\ell) = \lambda_s \cdot (wh)^{-1} \sum_{x \in \aleph} \|\nabla \delta_t^\ell(x)\|^2 + \lambda_r \cdot (wh)^{-1} \sum_{x \in \aleph} \|\delta_t^\ell(x)\|^2 \quad (\text{A.1})$$

with $\lambda_r, \lambda_s \in \{10^{-4}, 5 \times 10^{-4}, 10^{-3}, 5 \times 10^{-3}, 10^{-2}\}$. With reference to Eq. (3), in the implementation we have appropriate coefficients for the three terms:

$$\begin{aligned} L_{conj}^\ell = & \lambda_{cur}^\ell \mathcal{L}_c(\delta_t^\ell, f_{t-1}^\ell, f_t^\ell) \\ & + \lambda_{skip}^\ell \mathcal{L}_c(\delta_t^\ell, f_{t-1}^\ell, f_t^\ell) \\ & + \lambda_{low}^\ell \mathcal{L}_c(\delta_t^\ell, f_{t-1}^{\ell-1}, f_t^{\ell-1}) \\ & + \mathcal{R}(\delta_t^\ell) \end{aligned} \quad (\text{A.2})$$

with the abovementioned coefficients in the range $\{1, 2\} \times \{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$. Concerning the learning schedule T_{sched} (motion estimator warmup), we tried the following values $\{0, 2, 10\}$ laps, where 0 means no scheduling (the entire model is updated starting

¹⁰ Inspired to <https://github.com/usuyama/pytorch-unet> (MIT License), but with one fourth of filters in each layer. Also, the last layer has an output dimensionality of 32 features

¹¹ Available at <https://pytorch.org/vision/stable/models.html>.

¹² Available at <https://github.com/facebookresearch/moco>.

¹³ Available at <https://github.com/facebookresearch/moco-v3>.

¹⁴ Available at <https://github.com/zdaxie/PixPro>.

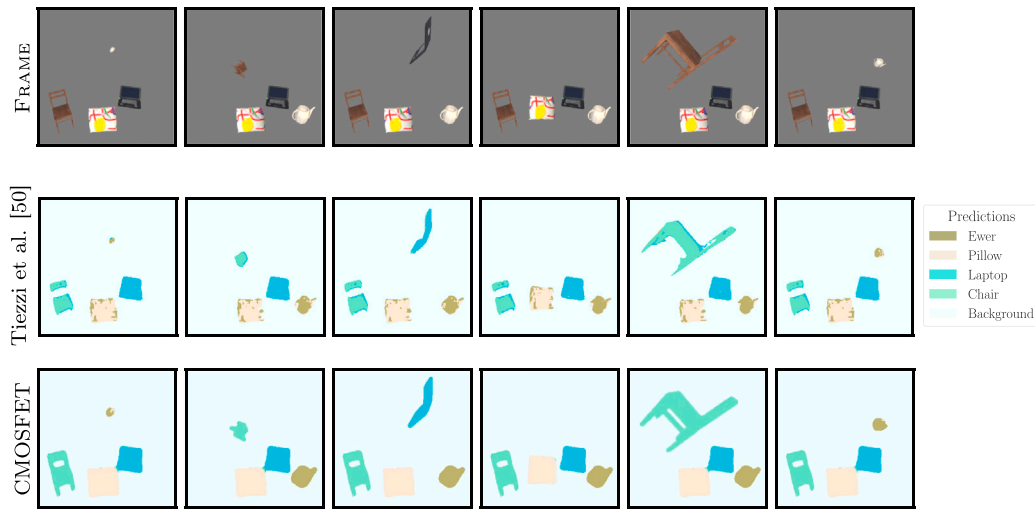


Fig. A.8. Six frames sampled from the EMPTYSPACE-rgb stream (first row). We compare the predictions yielded by the main competitor (second row) and the outputs by CMOSFET (third row). Different colors are used to denote different predictions.

Table A.2

Optimal parameters. The best selected hyperparameters drawn from the grids described in the text, for all the datasets. See the code for further details.

Parameters	EMPTYSPACE		SOLID	LIVINGROOM		RAT	HORSE
	BW	RGB	BW	BW	RGB	RGB	RGB
\aleph	0.7	1.0	1.0	1.0	1.0	0.7	1.0
α_f	10^{-3}	10^{-3}	10^{-2}	10^{-3}	10^{-3}	10^{-3}	10^{-4}
α_m	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-4}	5×10^{-4}	10^{-4}
η	100	100	100	200	200	200	200
λ_{low}^0	1	1	1	1	1	1	1
λ_{low}^1	0.1	1	1	1	1	0.01	0.01
λ_{cur}^0	10^{-4}	10^{-4}	10^{-2}	10^{-4}	10^{-4}	10^{-4}	10^{-2}
λ_{cur}^1	10^{-2}	10^{-4}	10^{-3}	10^{-2}	10^{-2}	10^{-4}	10^{-2}
λ_{skip}^0	0	0	0	0.2	0.2	2×10^{-5}	2×10^{-3}
λ_{skip}^1	0	0	0	0.2	0.2	2×10^{-5}	2×10^{-3}
τ	0.5	0.1	0.5	0.1	0.1	0.5	0.5
τ_m	1.5	1.5	2	0.7	0.7	1.3	1.3
τ_n	0	0	-0.3	-0.5	-0.5	0.5	0.5
τ_p	0.9	0.7	0.7	0.9	0.9	0.8	0.8
ξ	0.99	0.5	0.99	0.99	0.99	0.99	0.99
λ_p	10^{-3}	5×10^{-3}	10^{-3}	10^{-2}	10^{-2}	10^{-2}	10^{-2}
λ_s	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-3}	10^{-4}
T_{sched}	10	0	0	0	0	10	2

on the very first frame). In all our experiments, we reported the average results obtained over 10 runs, initializing random generators with seeds (`torch.manual_seed` in PyTorch) from the following list: {1234, 123 456, 12 345 678, 1 234 567 890, 9876, 98 765, 987 654, 9 876 543, 98 765 432, 987 654 321}.

We report in Table A.2 the best-found values for the hyperparameters.

A.3. Additional qualitative analysis

We report in Fig. A.8 some additional qualitative results of the predictions obtained leveraging the features by CMOSFET, on six frames from the EMPTYSPACE-rgb stream (first row). The model by Tiezzi, Marullo, Faggi et al. (2022) yields features which are not capable to correctly discriminate pixels from the chair legs/thinner segments, as long as some textures on the pillow: see for instance the borders of the pillow which are misled with the ewer (given the similar brightness/RGB values). Similarly, some parts of the ewer (the spout and the handle) are mistaken for pixels belonging to the pillow. Conversely,

CMOSFET allows the classification procedure (which is not-involved in the feature learning process) to almost completely disentangle all the different objects and their parts, even when they appear in peculiar poses and orientations (see the chair legs in the second and fifth image). As a general consideration, CMOSFET yields features (and thus predictions) which tend to slightly overflow with respect to the ground-truth object borders.

A.4. Code

We provide a Python implementation of CMOSFET, as well as instructions to run experiments.¹⁵ The README file contained in the archive describes how to execute the code, and it gives further details on the mapping between the notation used in the main paper for naming hyperparameters and the one used in the code (see “RUNNING EXPERIMENTS - Argument description/mapping with respect to the paper notation” therein).

A.5. Efficiency analysis

CMOSFET is independent of the specific structure of the model used for feature extraction and motion estimation. This allows users to decide on the most suitable balance between computational speed and prediction quality, ensuring real-time processing requirements are met if necessary. The model presented in the main paper can process video streams at approximately 15 fps on the machine detailed in the experimental section. However, when using the shared code, in which data logging is enabled (saving metrics, outputs, and other information to gather experimental evidence), the overall processing speed naturally decreases. On average, each run took between 40 and 60 min, depending on the stream being processed. With respect to the main competitor CMOSFET-Single, CMOSFET involves the additional cost of computing motion flows for the considered multiple levels, implying an increased computational expense.

Data availability

Link to code repository is provided in the paper.

¹⁵ <https://github.com/sailab-code/unsupervised-learning-feature-flow>.

References

- Aljundi, Rahaf, Kelchtermans, Klaas, & Tuytelaars, Tinne (2019). Task-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11254–11263).
- Arani, Elahe, Sarfraz, Fahad, & Zonooz, Bahram (2022). Learning fast, learning slow: A general continual learning method based on complementary learning system. In *International conference on learning representations*.
- Baker, Simon, Roth, Stefan, Scharstein, Daniel, Black, Michael J., Lewis, J. P., & Szeliski, Richard (2007). A database and evaluation methodology for optical flow. In *2007 IEEE 11th international conference on computer vision* (pp. 1–8).
- Betti, Alessandro, Boccignone, Giuseppe, Faggi, Lapo, Gori, Marco, & Melacci, Stefano (2021). Visual features and their own optical flow. *Frontiers in Artificial Intelligence*, 4.
- Betti, Alessandro, Faggi, Lapo, Gori, Marco, Tiezzi, Matteo, Marullo, Simone, Meloni, Enrico, et al. (2022). Continual learning through hamilton equations. In *conference on lifelong learning agents* (pp. 201–212). PMLR.
- Betti, Alessandro, Gori, Marco, & Melacci, Stefano (2022). *Deep learning to see-towards new foundations of computer vision*. Springer.
- Bregler, Christoph (1997). Learning and recognizing human dynamics in video sequences. In *proceedings of IEEE computer society conference on computer vision and pattern recognition* (pp. 568–574).
- Brox, Thomas, Bruhn, Andrés, Papenberger, Nils, & Weickert, Joachim (2004). High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision* (pp. 25–36). Springer.
- Cai, Zhaowei, Ravichandran, Avinash, Maji, Subhransu, Fowlkes, Charless, Tu, Zhouwen, & Soatto, Stefano (2021). Exponential moving average normalization for self-supervised and semi-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 194–203).
- Chen, Xinlei, Fan, Haoqi, Girshick, Ross, & He, Kaiming (2020). Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297.
- Chen, Xinlei, & He, Kaiming (2021). Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 15750–15758).
- Chen, Ting, Kornblith, Simon, Norouzi, Mohammad, & Hinton, Geoffrey (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607). PMLR.
- Chen, Liang-Chieh, Papandreou, George, Kokkinos, Iasonas, Murphy, Kevin, & Yuille, Alan L. (2017). Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4), 834–848.
- Chen*, Xinlei, Xie*, Saining, & He, Kaiming (2021). An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 9640–9649).
- De Lange, Matthias, Aljundi, Rahaf, Masana, Marc, Parisot, Sarah, Jia, Xu, Leonardis, Aleš, et al. (2021). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7), 3366–3385.
- Delange, Matthias, Aljundi, Rahaf, Masana, Marc, Parisot, Sarah, Jia, Xu, Leonardis, Ales, et al. (2021). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1.
- Dosovitskiy, Alexey, Fischer, Philipp, Ilg, Eddy, Häusser, Philip, Hazirbas, Caner, Golkov, Vladimir, et al. (2015). FlowNet: Learning optical flow with convolutional networks. In *2015 IEEE international conference on computer vision* (pp. 2758–2766).
- Geng, Chuanxing, Huang, Sheng-Jun, & Chen, Songcan (2020). Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1.
- Grill, Jean-Bastien, Strub, Florian, Althé, Florent, Tallec, Corentin, Richemond, Pierre, Buchatskaya, Elena, et al. (2020). Bootstrap your own latent-a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33, 21271–21284.
- Habibian, Amirhossein, Yahia, Haitam, Ben, Abati, Davide, Gavves, Efstratios, & Porikli, Fatih (2022). Delta distillation for efficient video processing. In *Proceedings of the 17th European conference on computer vision*. Springer.
- He, Kaiming, Fan, Haoqi, Wu, Yuxin, Xie, Saining, & Girshick, Ross (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729–9738).
- Horn, Berthold K. P., & Schunck, Brian. G. (1981). Determining optical flow. *Artificial Intelligence*, 17(1–3), 185–203.
- Jing, Li, Vincent, Pascal, LeCun, Yann, & Tian, Yuandong Understanding dimensional collapse in contrastive self-supervised learning, In: International conference on learning representations.
- Kingma, Diederik P., & Ba, Jimmy (2014). Dam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Le-Khac, Phuc H., Healy, Graham, & Smeaton, Alan F. (2020). Contrastive representation learning: A framework and review. *IEEE Access*, 8, 193907–193934.
- Li, Rui, Zhang, Yiheng, Qiu, Zhaofan, Yao, Ting, Liu, Dong, & Mei, Tao (2021). Motion-focused contrastive learning of video representations. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 2105–2114).
- Liang, Yongqing, Li, Xin, Jafari, Navid, & Chen, Jim (2020). Video object segmentation with adaptive feature bank and uncertain-region refinement. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), vol. 33, *Advances in neural information processing systems* (pp. 3430–3441). Curran Associates, Inc.
- Lin, Tsung-Yi, Maire, Michael, Belongie, Serge, Hays, James, Perona, Pietro, Ramanan, Deva, et al. (2014). Microsoft coco: Common objects in context. In *ECCV* (pp. 740–755). Springer.
- Liu, Xiao, Zhang, Fanjin, Hou, Zhenyu, Mian, Li, Wang, Zhaoyu, Zhang, Jing, et al. (2021). Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 35(1), 857–876.
- Madaan, Divyam, Yoon, Jaehong, Li, Yuanchun, Liu, Yunxin, & Hwang, Sung-Ju (2022). Representational continuity for unsupervised continual learning. In *International conference on learning representations*.
- Mahendran, Aravindh, Thewlis, James, & Vedaldi, Andrea (2019). Cross pixel optical-flow similarity for self-supervised learning. In *Computer vision—ACCV 2018: 14th Asian conference on computer vision, perth, Australia, December (2018) 2–6, revised selected papers, part v 14* (pp. 99–116). Springer.
- Marullo, Simone, Tiezzi, Matteo, Betti, Alessandro, Faggi, Lapo, Meloni, Enrico, & Melacci, Stefano (2022). Continual unsupervised learning for optical flow estimation with deep networks. In *Conference on lifelong learning agents* (pp. 183–200). PMLR.
- Meloni, Enrico, Betti, Alessandro, Faggi, Lapo, Marullo, Simone, Tiezzi, Matteo, & Melacci, Stefano (2021). Evaluating continual learning algorithms by generating 3d virtual environments. arXiv preprint arXiv:2109.07855.
- Meloni, Enrico, Pasqualini, Luca, Tiezzi, Matteo, Gori, Marco, & Melacci, Stefano (2020). Sailenv: Learning in virtual visual environments made simple. In *ICPR* (pp. 8906–8913).
- Ostrovsky, Yuri, Meyers, Ethan, Ganesh, Suma, Mathur, Umang, & Sinha, Pawan (2009). Visual parsing after recovery from blindness. *Psychological Science*, 20(12), 1484–1491.
- Pathak, Deepak, Girshick, Ross, Dollár, Piotr, Darrell, Trevor, & Hariharan, Bharath (2017). Learning features by watching objects move. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2701–2710).
- Qian, Rui, Meng, Tianjian, Gong, Boqing, Yang, Ming-Hsuan, Wang, Huisheng, Belongie, Serge, et al. (2021). Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6964–6974).
- Ranfll, René, Bochkovskiy, Alexey, & Koltun, Vladlen (2021). Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 12179–12188).
- Rao, Dushyant, Visin, Francesco, Rusu, Andrei, Pascanu, Razvan, Teh, Yee Whye, & Hadsell, Raia (2019). Continual unsupervised representation learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), vol. 32, *Advances in neural information processing systems*. Curran Associates, Inc.
- Ronneberger, Olaf, Fischer, Philipp, & Brox, Thomas (2015). U-net: Convolutional networks for biomedical image segmentation. In *MICCAI* (pp. 234–241). Springer.
- Scheirer, Walter J., Rezende Rocha, Anderson de, Sapkota, Archana, & Boulton, Terrance E. (2012). Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7), 1757–1772.
- Spelke, Elizabeth S. (1990). Principles of object perception. *Cognitive Science*, 14(1), 29–56.
- Sun, Lin, Jia, Kui, Chan, Tsung-Han, Fang, Yuqiang, Wang, Gang, & Yan, Shuicheng (2014). Df-sfa: Deeply-learned slow feature analysis for action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Sun, Deqing, Roth, Stefan, & Black, Michael J. (2010). Secrets of optical flow estimation and their principles. In *International conference on computer vision and pattern recognition* (pp. 2432–2439).
- Tarvainen, Antti, & Valpola, Harri (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), vol. 30, *Advances in neural information processing systems*. Curran Associates, Inc.
- Teed, Zachary, & Deng, Jia (2020). Raft: Recurrent all-pairs field transforms for optical flow.
- Tiezzi, Matteo, Marullo, Simone, Becattini, Federico, & Melacci, Stefano (2024). Continual neural computation. In *Joint European conference on machine learning and knowledge discovery in databases* (pp. 340–356). Springer.
- Tiezzi, Matteo, Marullo, Simone, Betti, Alessandro, Casoni, Michele, & Melacci, Stefano (2024). Bridging continual learning of motion and self-supervised representations. In *Proceedings of the 27th European conference on artificial intelligence*.
- Tiezzi, Matteo, Marullo, Simone, Betti, Alessandro, Meloni, Enrico, Faggi, Lapo, Gori, Marco, et al. (2022). Foveated neural computation. In *Joint European conference on machine learning and knowledge discovery in databases* (pp. 19–35). Springer.
- Tiezzi, Matteo, Marullo, Simone, Faggi, Lapo, Meloni, Enrico, Betti, Alessandro, & Melacci, Stefano (2022). Stochastic coherence over attention trajectory for continuous learning in video streams. In *Proceedings of the thirty-first international joint conference on artificial intelligence, IJCAI-22* (pp. 3480–3486).
- Tiezzi, Matteo, Melacci, Stefano, Betti, Alessandro, Maggini, Marco, & Gori, Marco (2020). Focus of attention improves information transfer in visual features. In *NeurIPS 2020-34th conference on neural information processing systems*.

- van de Ven, Gido M., Tuytelaars, Tinne, & Tolias, Andreas S. (2022). Three types of incremental learning. *Nature Machine Intelligence*, 4(12), 1185–1197.
- Wertheimer, Max (1938). Laws of organization in perceptual forms. In *A source book of gestalt psychology* (pp. 71–88).
- Wiskott, Laurenz, & Sejnowski, Terrence J (2002). Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4), 715–770.
- Xie, Zhenda, Lin, Yutong, Zhang, Zheng, Cao, Yue, Lin, Stephen, & Hu, Han (2021). Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 16684–16693).
- Xiong, Y, Ren, M, Zeng, W, & Waabi, R (2021). Self-supervised representation learning from flow equivariance. In *Proceedings of the IEEE international conference on computer vision* (pp. 10171–10180).
- Zanca, Dario, Gori, Marco, Melacci, Stefano, & Rufa, Alessandra (2020). Gravitational models explain shifts on human visual attention. *Scientific Reports*, 10(1), 1–9.
- Zhai, Mingliang, Xiang, Xuezhai, Lv, Ning, & Kong, Xiangdong (2021). Optical flow and scene flow estimation: A survey. *Pattern Recognition*, 114, Article 107861.
- Zhang, Chaoning, Zhang, Kang, Zhang, Chenshuang, Pham, X, Yoo, Chang D, & Kweon, In So How does simsiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning. In *International conference on learning representations*.
- Zhang, Chaoning, Zhang, Kang, Zhang, Chenshuang, Pham, X, Yoo, Chang D., & Kweon, In So (2022). How does simsiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning. In *International conference on learning representations*.
- Zhou, Bolei, Zhao, Hang, Puig, Xavier, Xiao, Tete, Fidler, Sanja, Barriuso, Adela, et al. (2019). Semantic understanding of scenes through the ade20k dataset. *IJCV*, 127(3), 302–321.