BALANCING DIFFERENTIAL DISCRIMINATIVE KNOWLEDGE FOR CLOTH-IRRELEVANT LIFELONG PERSON RE-IDENTIFICATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Lifelong person re-identification (L-ReID) focuses on learning sequentially collected datasets from different domains to match the same person. Advanced L-ReID methods typically balance the domain gap between different datasets via domain knowledge modeling, such as knowledge rectification or distribution prototyping. However, existing methods dismiss balancing discriminative knowledge within different datasets, resulting in conflicts when sequentially accumulating differential discriminative information in different datasets, e.g., sequentially learning cloth-changing/cloth-consistent knowledge simultaneously, which brings critical catastrophic forgetting problems of old discriminative knowledge. In this paper, we focus on a new but practical task called Cloth-Irrelevant Lifelong Person Re-identification (CIL-ReID), which requires matching the same person wearing different clothing using sequentially collected data. To tackle the above issue, we proposed an Adaptive Discriminative Knowledge Consolidation (ADKC) framework to balance the discriminative information of different domains on L-ReID. Specifically, we propose a Selective Knowledge Forgetting (SKF) module to correct potential overfitting to specific discrimination (e.g., clothing information) based on new knowledge. In addition, we design a Selective Knowledge Retention(SKR) module to adaptively compensate for the potential lack of discriminative information based on old knowledge and accelerate differential discrimination into a unified framework. To validate our method, two CCL-ReID benchmarks are first established, while extensive experiments on the above two benchmark datasets demonstrate that our method leads to existing advanced methods in the CCL-ReID task.

1 INTRODUCTION

As a fundamental task in computer vision, person re-identification (ReID) focuses on retrieving the same person across different locations. However, recent studies have revealed its limitations when it comes to involving continual learning of training data, which is also called the catastrophic forgetting problem. To tackle the above issue, lifelong person re-identification (L-ReID) intends to continuously accumulate knowledge from old and new training data, balancing preserving old knowledge while acquiring new knowledge.

To this end, some L-ReID methods preserve old knowledge by storing and replaying additional historical data when training on new data. However, despite some progress, these methods often exhibit a significant performance decline in the data privacy scenario that prohibits data replay. Therefore, the other methods are devoted to preserving old knowledge without historical data, known as the non-replay L-ReID method. Specifically, acknowledging the domain discrepancies among continuously collected datasets, e.g., style, colour, and resolution variations, these methods often balance old and new knowledge by developing robust modelling strategies, including distribution modelling, image patch modelling, and relationship modelling, to mitigate the discrepancies across different data domains.

However, these methods typically overlook the discrepancies of identical information among different datasets, resulting in challenges in effectively accumulating various knowledge across do-

mains. Specifically, as shown in Figure 1, existing L-ReID methods will inevitably handle both cloth-consistent and cloth-changing data when continuously collecting new data. Notably, while clothing information may often effective in the cloth-consistent scenario, it does not always perform equally in the cloth-changing one, which relies more on other discriminative information, such the figure, body, and shape.

Consequently, these methods often encountered significant challenges when adapting to new scenarios and preserving the performance in old scenarios, requiring balancing the individual discriminations and the domain discrepancies across datasets. However, drastic changes in the sequentially collected data bring serious conflicts when matching the same person, ultimately hindering the average performance of the both cloth-consistent and cloth-changing scenarios.

Inspired by the above observations, in this paper, we focus on a practical but challenging task, called Cloth-Irrelevant Lifelong person Re-IDentification (CIL-ReID), which requires employing streaming cloth-consistent and cloth-changing data to perform lifelong learning and perform well in both scenarios. To this end, we proposed an Adaptive Discriminative Knowledge Consolidation (ADKC) framework to balance the discriminative information of different domains on L-ReID. Specifically, we propose a Selective Knowledge Forgetting (SKF) module to correct potential overfitting to specific discrimination (e.g., clothing information) based on new knowledge. In addition, we design a Selective Knowledge Retention(SKR) module to adaptively compensate for the potential lack of discriminative information based on old knowledge and accelerate differential discrimination into a unified framework.

Method	Cloth-Consistent										Average		Cloth-Changing				Average	
	Market-1501		LTCC		PRCC		MSMT17-V2		CUHK03				LTCC		PRCC			
	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1
JointTrain	64.1	82.5	42.6	62.1	94.6	98.7	18.4	40.8	44.4	46.4	52.8	66.1	10.1	23.0	32.7	33.8	21.4	28.4
SFT	28.5	52.0	28.5	49.3	92.5	97.3	7.0	19.6	44.0	45.6	40.1	52.8	6.9	15.3	21.8	21.2	14.4	18.3
LwF ?	44.5	65.8	21.6	40.0	87.4	91.3	4.0	11.6	25.5	25.0	36.6	46.7	5.9	12.5	25.9	26.7	15.9	19.6
AKA ?	48.0	69.5	25.4	45.1	88.1	93.3	4.2	12.0	31.2	31.2	39.4	50.2	6.5	12.8	26.5	26.7	16.5	19.8
PatchKD ?	**68.0**	**85.5**	30.8	54.7	93.5	96.5	5.7	15.6	33.2	32.9	46.2	57.0	7.2	17.9	26.1	26.0	16.7	22.0
LSTKC	39.9	63.4	39.6	60.4	95.9	98.9	11.5	29.2	**48.1**	**50.1**	47.0	60.4	8.3	19.4	24.0	22.9	16.2	21.2
DKP	36.1	60.7	39.9	59.8	96.7	98.8	12.8	30.6	34.9	36.7	44.1	57.3	**9.6**	22.4	**34.6**	32.8	**22.1**	27.6
Ours	51.6	74.8	**45.9**	**65.3**	**98.4**	**99.5**	**17.7**	**39.8**	36.9	38.4	**50.1**	**63.6**	**9.6**	**26.3**	33.7	**34.3**	21.7	**30.3**

Table 1: Performance on training *Order-1*: Market-1501→LTCC→PRCC→MSMT17-V2→CUHK03.

1.1 STYLE

Papers to be submitted to ICLR 2025 must be prepared according to the instructions presented here.

Authors are required to use the ICLR LaTeX style files obtainable at the ICLR website. Please make sure you use the current files and not previous versions. Tweaking the style files may be grounds for rejection.

1.2 RETRIEVAL OF STYLE FILES

The style files for ICLR and other conference information are available online at:

`http://www.iclr.cc/`

The file `iclr2025_conference.pdf` contains these instructions and illustrates the various formatting requirements your ICLR paper must satisfy. Submissions must be made using LaTeX and the style files `iclr2025_conference.sty` and `iclr2025_conference.bst` (to be used with LaTeX2e). The file `iclr2025_conference.tex` may be used as a "shell" for writing your paper. All you have to do is replace the author, title, abstract, and text of the paper with your own.

The formatting instructions contained in these style files are summarized in sections 2, 3, and 4 below.

2 GENERAL FORMATTING INSTRUCTIONS

The text must be confined within a rectangle 5.5 inches (33 picas) wide and 9 inches (54 picas) long. The left margin is 1.5 inch (9 picas). Use 10 point type with a vertical spacing of 11 points. Times New Roman is the preferred typeface throughout. Paragraphs are separated by 1/2 line space, with no indentation.

Paper title is 17 point, in small caps and left-aligned. All pages should start at 1 inch (6 picas) from the top of the page.

Authors' names are set in boldface, and each name is placed above its corresponding address. The lead author's name is to be listed first, and the co-authors' names are set to follow. Authors sharing the same address can be on the same line.

Please pay special attention to the instructions in section 4 regarding figures, tables, acknowledgments, and references.

There will be a strict upper limit of 10 pages for the main text of the initial submission, with unlimited additional pages for citations.

3 HEADINGS: FIRST LEVEL

First level headings are in small caps, flush left and in point size 12. One line space before the first level heading and 1/2 line space after the first level heading.

3.1 HEADINGS: SECOND LEVEL

Second level headings are in small caps, flush left and in point size 10. One line space before the second level heading and 1/2 line space after the second level heading.

3.1.1 HEADINGS: THIRD LEVEL

Third level headings are in small caps, flush left and in point size 10. One line space before the third level heading and 1/2 line space after the third level heading.

4 CITATIONS, FIGURES, TABLES, REFERENCES

These instructions apply to everyone, regardless of the formatter being used.

4.1 CITATIONS WITHIN THE TEXT

Citations within the text should be based on the `natbib` package and include the authors' last names and year (with the "et al." construct for more than two authors). When the authors or the publication are included in the sentence, the citation should not be in parenthesis using `\citet{}` (as in "See Hinton et al. (2006) for more information."). Otherwise, the citation should be in parenthesis using `\citep{}` (as in "Deep learning shows promise to make progress towards AI (Bengio & LeCun, 2007).").

The corresponding references are to be listed in alphabetical order of authors, in the REFERENCES section. As to the format of the references themselves, any style is acceptable as long as it is used consistently.

4.2 FOOTNOTES

Indicate footnotes with a number[1] in the text. Place the footnotes at the bottom of the page on which they appear. Precede the footnote with a horizontal rule of 2 inches (12 picas).[2]

[1] Sample of the first footnote

[2] Sample of the second footnote

Table 2: Sample table title

PART	DESCRIPTION
Dendrite	Input terminal
Axon	Output terminal
Soma	Cell body (contains cell nucleus)

4.3 FIGURES

All artwork must be neat, clean, and legible. Lines should be dark enough for purposes of reproduction; art work should not be hand-drawn. The figure number and caption always appear after the figure. Place one line space before the figure caption, and one line space after the figure. The figure caption is lower case (except for first word and proper nouns); figures are numbered consecutively.

Make sure the figure caption does not get separated from the figure. Leave sufficient space to avoid splitting the figure and figure caption.

You may use color figures. However, it is best for the figure captions and the paper body to make sense if the paper is printed either in black/white or in color.

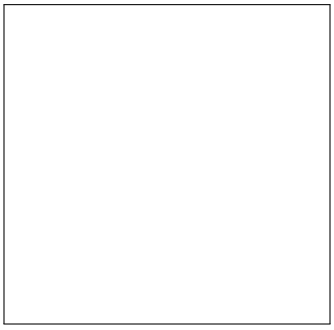

Figure 1: Sample figure caption.

4.4 TABLES

All tables must be centered, neat, clean and legible. Do not use hand-drawn tables. The table number and title always appear before the table. See Table 2.

Place one line space before the table title, one line space after the table title, and one line space after the table. The table title must be lower case (except for first word and proper nouns); tables are numbered consecutively.

5 DEFAULT NOTATION

In an attempt to encourage standardized notation, we have included the notation file from the textbook, *Deep Learning* Goodfellow et al. (2016) available at `https://github.com/goodfeli/dlbook_notation/`. Use of this style is not required and can be disabled by commenting out `math_commands.tex`.

Numbers and Arrays

a	A scalar (integer or real)
$\boldsymbol{a}$	A vector
$\boldsymbol{A}$	A matrix
$\mathbf{A}$	A tensor
$\boldsymbol{I}_n$	Identity matrix with n rows and n columns
$\boldsymbol{I}$	Identity matrix with dimensionality implied by context
$\boldsymbol{e}^{(i)}$	Standard basis vector $[0, \dots, 0, 1, 0, \dots, 0]$ with a 1 at position i
$\text{diag}(\boldsymbol{a})$	A square, diagonal matrix with diagonal entries given by $\boldsymbol{a}$
a	A scalar random variable
$\mathbf{a}$	A vector-valued random variable
$\mathbf{A}$	A matrix-valued random variable

Sets and Graphs

$\mathbb{A}$	A set
$\mathbb{R}$	The set of real numbers
$\{0, 1\}$	The set containing 0 and 1
$\{0, 1, \dots, n\}$	The set of all integers between 0 and n
$[a, b]$	The real interval including a and b
$(a, b]$	The real interval excluding a but including b
$\mathbb{A} \backslash \mathbb{B}$	Set subtraction, i.e., the set containing the elements of $\mathbb{A}$ that are not in $\mathbb{B}$
$\mathcal{G}$	A graph
$Pa_{\mathcal{G}}(\mathrm{x}_i)$	The parents of x_i in $\mathcal{G}$

Indexing

a_i	Element i of vector $\boldsymbol{a}$, with indexing starting at 1
a_{-i}	All elements of vector $\boldsymbol{a}$ except for element i
$A_{i,j}$	Element i, j of matrix $\boldsymbol{A}$
$\boldsymbol{A}_{i,:}$	Row i of matrix $\boldsymbol{A}$
$\boldsymbol{A}_{:,i}$	Column i of matrix $\boldsymbol{A}$
$\boldsymbol{A}_{i,j,k}$	Element (i, j, k) of a 3-D tensor $\mathbf{A}$
$\mathbf{A}_{:,:,i}$	2-D slice of a 3-D tensor
a_i	Element i of the random vector $\mathbf{a}$

Calculus

$\frac{dy}{dx}$	Derivative of y with respect to x
$\frac{\partial y}{\partial x}$	Partial derivative of y with respect to x
$\nabla_{\boldsymbol{x}} y$	Gradient of y with respect to $\boldsymbol{x}$
$\nabla_{\boldsymbol{X}} y$	Matrix derivatives of y with respect to $\boldsymbol{X}$
$\nabla_{\mathbf{X}} y$	Tensor containing derivatives of y with respect to $\mathbf{X}$
$\frac{\partial f}{\partial \boldsymbol{x}}$	Jacobian matrix $\boldsymbol{J} \in \mathbb{R}^{m \times n}$ of $f : \mathbb{R}^n \to \mathbb{R}^m$
$\nabla^2_{\boldsymbol{x}} f(\boldsymbol{x})$ or $\boldsymbol{H}(f)(\boldsymbol{x})$	The Hessian matrix of f at input point $\boldsymbol{x}$
$\int f(\boldsymbol{x}) d\boldsymbol{x}$	Definite integral over the entire domain of $\boldsymbol{x}$
$\int_{\mathbb{S}} f(\boldsymbol{x}) d\boldsymbol{x}$	Definite integral with respect to $\boldsymbol{x}$ over the set $\mathbb{S}$

Probability and Information Theory

$P(\mathrm{a})$	A probability distribution over a discrete variable
$p(\mathrm{a})$	A probability distribution over a continuous variable, or over a variable whose type has not been specified
$\mathrm{a} \sim P$	Random variable a has distribution P
$\mathbb{E}_{\mathrm{x} \sim P}[f(x)]$ or $\mathbb{E} f(x)$	Expectation of $f(x)$ with respect to $P(\mathrm{x})$
$\mathrm{Var}(f(x))$	Variance of $f(x)$ under $P(\mathrm{x})$
$\mathrm{Cov}(f(x), g(x))$	Covariance of $f(x)$ and $g(x)$ under $P(\mathrm{x})$
$H(\mathrm{x})$	Shannon entropy of the random variable x
$D_{\mathrm{KL}}(P \Vert Q)$	Kullback-Leibler divergence of P and Q
$\mathcal{N}(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution over $\boldsymbol{x}$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$

Functions

$f : \mathbb{A} \to \mathbb{B}$	The function f with domain $\mathbb{A}$ and range $\mathbb{B}$
$f \circ g$	Composition of the functions f and g
$f(\boldsymbol{x}; \boldsymbol{\theta})$	A function of $\boldsymbol{x}$ parametrized by $\boldsymbol{\theta}$. (Sometimes we write $f(\boldsymbol{x})$ and omit the argument $\boldsymbol{\theta}$ to lighten notation)
$\log x$	Natural logarithm of x
$\sigma(x)$	Logistic sigmoid, $\frac{1}{1 + \exp(-x)}$
$\zeta(x)$	Softplus, $\log(1 + \exp(x))$
$\Vert \boldsymbol{x} \Vert_p$	L^p norm of $\boldsymbol{x}$
$\Vert \boldsymbol{x} \Vert$	L^2 norm of $\boldsymbol{x}$
x^+	Positive part of x, i.e., $\max(0, x)$
$\mathbf{1}_{\mathrm{condition}}$	is 1 if the condition is true, 0 otherwise

6 FINAL INSTRUCTIONS

Do not change any aspects of the formatting parameters in the style files. In particular, do not modify the width or length of the rectangle the text should fit into, and do not change font sizes (except perhaps in the REFERENCES section; see below). Please note that pages should be numbered.

7 PREPARING POSTSCRIPT OR PDF FILES

Please prepare PostScript or PDF files with paper size "US Letter", and not, for example, "A4". The -t letter option on dvips will produce US Letter files.

Consider directly generating PDF files using `pdflatex` (especially if you are a MiKTeX user). PDF figures must be substituted for EPS figures, however.

Otherwise, please generate your PostScript and PDF files with the following commands:

```
dvips mypaper.dvi -t letter -Ppdf -G0 -o mypaper.ps
ps2pdf mypaper.ps mypaper.pdf
```

7.1 MARGINS IN LATEX

Most of the margin problems come from figures positioned by hand using `\special` or other commands. We suggest using the command `\includegraphics` from the graphicx package. Always specify the figure width as a multiple of the line width as in the example below using .eps graphics

```
\usepackage[dvips]{graphicx} ...
\includegraphics[width=0.8\linewidth]{myfile.eps}
```

or

```
\usepackage[pdftex]{graphicx} ...
\includegraphics[width=0.8\linewidth]{myfile.pdf}
```

for .pdf graphics. See section 4.4 in the graphics bundle documentation (`http://www.ctan.org/tex-archive/macros/latex/required/graphics/grfguide.ps`)

A number of width problems arise when LaTeX cannot properly hyphenate a line. Please give LaTeX hyphenation hints using the `\-` command.

AUTHOR CONTRIBUTIONS

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors.

ACKNOWLEDGMENTS

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper.

REFERENCES

Yoshua Bengio and Yann LeCun. Scaling learning algorithms towards AI. In *Large Scale Kernel Machines*. MIT Press, 2007.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.

Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.

A APPENDIX

You may include other additional sections here.