

One-shot Joint Extraction, Registration and Segmentation of Neuroimaging Data

Yao Su
Worcester Polytechnic Institute
ysu6@wpi.edu

Zhentian Qian
Worcester Polytechnic Institute
zqian@wpi.edu

Lei Ma
Worcester Polytechnic Institute
lma5@wpi.edu

Lifang He
Lehigh University
lih319@lehigh.edu

Xiangnan Kong
Worcester Polytechnic Institute
xkong@wpi.edu

ABSTRACT

Brain extraction, registration and segmentation are indispensable preprocessing steps in neuroimaging studies. The aim is to extract the brain from raw imaging scans (*i.e.*, extraction step), align it with a target brain image (*i.e.*, registration step) and label the anatomical brain regions (*i.e.*, segmentation step). Conventional studies typically focus on developing separate methods for the extraction, registration and segmentation tasks in a supervised setting. The performance of these methods is largely contingent on the quantity of training samples and the extent of visual inspections carried out by experts for error correction. Nevertheless, collecting voxel-level labels and performing manual quality control on high-dimensional neuroimages (*e.g.*, 3D MRI) are expensive and time-consuming in many medical studies. In this paper, we study the problem of one-shot joint extraction, registration and segmentation in neuroimaging data, which exploits only one labeled template image (*a.k.a.* atlas) and a few unlabeled raw images for training. We propose a unified end-to-end framework, called JERS, to jointly optimize the extraction, registration and segmentation tasks, allowing feedback among them. Specifically, we use a group of extraction, registration and segmentation modules to learn the extraction mask, transformation and segmentation mask, where modules are interconnected and mutually reinforced by self-supervision. Empirical results on real-world datasets demonstrate that our proposed method performs exceptionally in the extraction, registration and segmentation tasks.

CCS CONCEPTS

• Information systems → Data mining; • Computing methodologies → 3D imaging; Matching; Image segmentation.

KEYWORDS

brain extraction, skull stripping, registration, alignment, segmentation, joint learning, one-shot

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0103-0/23/08...\$15.00
<https://doi.org/10.1145/3580305.3599452>




Figure 1: The problem of one-shot joint extraction, registration and segmentation in neuroimaging data. Given a set of unlabeled raw images of the patients' heads, a standard template image of the brain and the segmentation labels of the template image, the goal is to train a model to perform brain extraction, registration and segmentation tasks simultaneously on raw images of new patients' heads.

ACM Reference Format:

Yao Su, Zhentian Qian, Lei Ma, Lifang He, and Xiangnan Kong. 2023. One-shot Joint Extraction, Registration and Segmentation of Neuroimaging Data. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3580305.3599452>

1 INTRODUCTION

Background. Brain extraction (*a.k.a.* skull stripping), registration and segmentation serve as preliminary yet indispensable steps in many neuroimaging studies, such as anatomical and functional analysis [4, 33, 50, 52], brain networks discovering [11, 24, 27, 28, 53, 54], multi-modality fusion [6, 26], diagnostic assistance [18, 46], and brain region studies [8, 25]. The brain extraction targets the removal of non-cerebral tissues (*e.g.*, skull, dura, and scalp) from a patient's head scan; the registration step aims to align the extracted brain with a standard brain template; the segmentation step intends to label the anatomical brain regions in the raw imaging scan. These three tasks serve as crucial preprocessing steps in many neuroimaging studies. For example, in brain functional and anatomical analysis, upon extracting and aligning the brain, the interference of non-cerebral tissues, imaging modalities, and viewpoints can be eliminated, thereby enabling accurate quantification of shifts in the shape, size, and signal; and labeled anatomical brain regions (*e.g.*,

frontal lobe, cerebellum, etc.) can be used to guide the structural diagnosis. In Alzheimer’s disease diagnosis, the brain across subjects needs to be first extracted from raw brain imaging scans and then aligned with a standard template to counteract inter-individual variations and perform brain function analysis (e.g., discovering the brain network connectivity). Meanwhile, the intra-individual structural lesions (e.g., brain atrophy) across different pathological stages need to be monitored in anatomical analysis (e.g., identify the corresponding anatomical brain region and measure its alteration of brain volume). These essential processing steps help doctors make a comprehensive and accurate diagnosis.

State-of-the-Art. The literature extensively explores brain extraction, registration, and segmentation problems [1, 7, 12, 22, 23, 29, 42]. Conventional approaches primarily emphasize the development of separate methods for extraction [23, 29], registration [12, 42], and segmentation [1, 7, 22] under supervised settings. However, within the domain of medical studies, the process of obtaining annotations for brain location, image transformations, and segmentation is often accompanied by significant expenses, necessitating expertise, and substantial time, especially when dealing with high-dimensional neuroimages (e.g., 3D MRI). To overcome this limitation, recent works [5, 10, 38, 39, 41, 56] introduce a three-step approach for one-shot extraction, registration and segmentation by using automated brain extraction tools [10, 38, 39, 41], unsupervised registration and segmentation models with direct warping [5, 20, 56], as shown in Figure 2(a). However, these approaches often rely on manual quality control to correct intermediate results before performing subsequent tasks, which is time-consuming, labor-intensive, and subject to variability, thus hampering overall efficiency and performance. More recently, joint extraction-registration method [45] and joint registration-segmentation methods [16, 35, 51] are introduced to solve the problem in a two-stage design, as shown in Figure 2(b) and Figure 2(c). However, partial joint learning neglects the potential relationship among all tasks and negatively impacts overall performance.

Problem Definition. This paper investigates the problem of one-shot joint brain extraction, registration, and segmentation, as shown in Figure 1. The goal is to capture the connections among three tasks to mutually boost their performance in a one-shot training scenario. Notably, the extraction, registration and segmentation labels of the raw image are not available. We expect to perform the three tasks simultaneously with only one labeled template.

Challenges. Despite its value and significance, the problem of one-shot joint extraction, registration and segmentation has not been studied before and is very challenging due to its unique characteristics listed below:

- *Lack of labels for extraction:* Traditional learning-based extraction methods rely on a substantial number of training samples with accurate ground truth labels. However, collecting voxel-level labels for high-dimensional neuroimaging data is a resource-intensive and time-consuming endeavor.
- *Lack of labels for registration:* Obtaining the accurate ground truth transformation between raw and template images poses significant challenges. While unsupervised registration methods [5, 56] optimize transformation parameters by maximizing image similarity, their effectiveness is contingent upon the prior removal of




Figure 2: Related works in one-shot brain extraction, registration and segmentation.

non-brain tissue from the raw image. Failing to do this may lead to erroneous transformations, rendering the registration invalid.

- *Lack of labels for segmentation:* Collecting the voxel-level segmentation labels is also difficult. Although we provide a template with its segmentation labeled (in template image space), the segmentation (in raw image space) of the raw image is not available.

- *Dependencies among extraction, registration and segmentation:* Conventional research typically treats extraction, registration, and segmentation tasks separately. However, these tasks exhibit a high degree of interdependency. The accuracy of registration and segmentation tasks heavily relies on the extraction task. The registration process assists the extraction task in capturing cerebral/non-cerebral information from raw and template images, and providing segmentation labels for guiding the segmentation task. The segmentation task can inversely force the extraction and registration tasks to provide precise results. Thus, a holistic solution is required to effectively manage the interdependencies among these tasks.

Proposed Method. To address the aforementioned challenges, we propose JERS, a unified end-to-end framework for joint brain extraction, registration, and segmentation. Figure 2 showcases a comparison between our method and state-of-the-art approaches. Specifically, JERS contains a group of extraction, registration and segmentation modules, where the extraction module gradually eliminates the non-brain tissue from the raw image, producing an extracted brain image; the registration module incrementally aligns the extracted image with the template and warps the template’s

segmentation label in the raw image space to guide the segmentation module; the segmentation module generates a segmentation label for the raw image and provides feedback to extraction and registration modules. These three modules help each other to boost extraction, registration and segmentation performance simultaneously. By bridging these three modules end-to-end, we achieve a joint optimization with only one labeled template assistance.

Upon evaluation on public brain MRI datasets, our proposed method significantly outperforms state-of-the-art techniques in extraction, registration, and segmentation accuracy.

2 PRELIMINARIES

In this section, we begin by introducing the relevant concepts and notations. We then proceed to formally define the problem of one-shot joint brain extraction, registration, and segmentation.

2.1 Notations and Definitions

Definition 1 (Source, target and target segmentation mask). Suppose we are given a training dataset $\mathcal{D} = \{\{S_i\}_{i=1}^Z, (T, B)\}$ that consists of Z source images $S_i \in \mathbb{R}^{W \times H \times D}$, and a pair of target image $T \in \mathbb{R}^{W \times H \times D}$ and its corresponding segmentation mask $B \in \{0, 1\}^{C \times W \times H \times D}$. Here, the source image S_i is the raw MRI scan of a patient's head, the target T is a standard template of the brain, and the segmentation mask B is the one-hot encoding of the target T segmentation. W , H , and D denote the width, height and depth dimensions of the 3D images, C denotes the number of anatomical labels (e.g., the number of labelled brain regions). To ensure simplicity, we make the assumption that the source and target images are resized to the same dimension, denoted as $W \times H \times D$. Next, we omit the subscript i of S_i for ease of notation.

Definition 2 (Brain extraction mask). Brain extraction mask $M \in \{0, 1\}^{W \times H \times D}$ is a binary tensor of identical dimensions to the source image S . It represents cerebral tissues in S with a value of 1 and non-cerebral tissues with 0. The extracted image $E = S \circ M$ is obtained by applying the M on S via a element-wise product $\circ$.

Definition 3 (Affine transformation and warped image). In order to maintain generality, we consider the transformation in the registration task to be affine-based. Extending this work to encompass other types of registration, such as nonlinear/deformable registration, is straightforward. The affine transformation parameters $\mathbf{a} \in \mathbb{R}^{12}$ is a vector used to parameterized an 3D affine transformation matrix $\mathbf{A} \in \mathbb{R}^{4 \times 4}$. The warped image $\mathbf{W} = \mathcal{T}(\mathbf{E}, \mathbf{a})$ results from applying the affine transformation on the extracted image $\mathbf{E}$, where $\mathcal{T}(\cdot, \cdot)$ denotes the affine transformation operator. At the voxel level, the relationship between $\mathbf{W}$ and $\mathbf{E}$ can be expressed as:

$$\mathbf{W}_{xyz} = \mathbf{E}_{x'y'z'}, \quad (1)$$

where the correspondences between coordinates x, y, z and x', y', z' are determined on the affine transformation matrix $\mathbf{A}$:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \mathbf{A} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} a_1 & a_2 & a_3 & a_4 \\ a_5 & a_6 & a_7 & a_8 \\ a_9 & a_{10} & a_{11} & a_{12} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}. \quad (2)$$

Definition 4 (Source segmentation mask). Brain segmentation mask $\mathbf{R} \in \{0, 1\}^{C \times W \times H \times D}$ is a binary tensor with the first dimension being the number of anatomical labels and the rest dimensions




Figure 3: A demonstration of extraction, registration and segmentation functions.

identical to the source image S . Every point in the source image S is densely labeled according to its anatomical structure, encoded in a one-hot vector at the corresponding coordinate of $\mathbf{R}$.

2.2 Problem Formulation

The goal of joint brain extraction, registration, and segmentation is to collectively learn the extraction function $f_\theta : \mathbb{R}^{W \times H \times D} \rightarrow \mathbb{R}^{W \times H \times D}$, the registration function $g_\phi : \mathbb{R}^{W \times H \times D} \times \mathbb{R}^{W \times H \times D} \rightarrow \mathbb{R}^{12}$, and the segmentation function $h_\psi : \mathbb{R}^{W \times H \times D} \rightarrow \mathbb{R}^{C \times W \times H \times D}$, as shown in Figure 3. Specifically, the extraction function $f_\theta(\cdot)$ utilizes the source image S as input to predicts a brain extraction mask $\hat{M} = f_\theta(S)$. The registration function $g_\phi(\cdot, \cdot)$ takes the extracted brain image $\hat{E} = \hat{M} \circ S$ and the target image T to predict the affine transformation parameter $\hat{\mathbf{a}} = g_\phi(\hat{E}, T)$, then obtaining the warped image $\hat{\mathbf{W}} = \mathcal{T}(\hat{E}, \hat{\mathbf{a}})$. The warped segmentation mask $\hat{\mathbf{V}} = \mathcal{T}(B, \hat{\mathbf{a}}^{-1})$ is generated by warping the target segmentation mask B using inversed affine transformation $\hat{\mathbf{a}}^{-1}$. Finally, the segmentation function $h_\psi(\cdot)$ takes the source image S as input to predict a source brain segmentation mask $\hat{\mathbf{R}} = h_\psi(S)$. The optimal parameter θ^*, ϕ^* and ψ^* can be found by solving the following optimization problem:

$$\theta^*, \phi^*, \psi^* = \arg \min_{\theta, \phi, \psi} \sum_{(S, T, B) \in \mathcal{D}} \left[\mathcal{L}_{sim}(\hat{\mathbf{W}}, T) + \lambda \mathcal{L}_{seg}(\hat{\mathbf{R}}, \hat{\mathbf{V}}) \right], \quad (3)$$

where the image pair (S, T, B) is sampled from the training dataset $\mathcal{D}$. $\mathcal{L}_{sim}(\cdot, \cdot)$ is image dissimilarity criteria, e.g., mean square error and $\mathcal{L}_{seg}(\cdot, \cdot)$ is segmentation dissimilarity criteria, e.g., cross entropy error. These two criteria guide a joint optimization of extraction, registration and segmentation functions, allowing feedback among them.

To the best of our knowledge, this work is the first endeavor in finding an optimal solution for the one-shot joint brain image extraction, registration, and segmentation problem. Our approach eliminates the necessity of labeling the brain extraction masks, transformation, and segmentation masks of the source image. We only require one pair of a target image and its corresponding segmentation mask to guide the training, as opposed to other fully supervised methods [1, 7, 12, 22, 23, 29, 42].




Figure 4: An overview of proposed JERS. **Extraction module** takes the raw source image S as input, and gradually produces the extracted brain image E^M after M stages of extraction. The final extracted image is E^M . **Registration module** takes the extracted brain image E^M , target image T and target segmentation mask B as inputs, and incrementally aligns E^M with T through N stages of registration. Then, it generates the warped segmentation mask V by inversely transforming the target segmentation mask B . **Segmentation module** takes the raw source image S as input, and output the brain segmentation mask R . Two loss terms couple all modules together, allowing them to perform joint learning. The final output of JERS is the extracted brain image E^M constituting only cerebral tissues, the warped image W^N aligning with the target image T , and the brain segmentation mask R indicating anatomical regions of the source image S .

3 OUR APPROACH

Overview. Figure 4 presents an overview of the proposed JERS framework for the one-shot joint brain extraction, registration, and segmentation problem. Our method is an end-to-end deep neural network consisting of three main branches: 1) *Extraction Module* takes the raw source image S as input, and progressively produces the extracted brain image E^M after M stages of extraction; 2) *Registration Module* takes the extracted brain image E^M , target image T and target segmentation mask B as inputs, and incrementally aligns E^M with T through N stages of registration. Then, it generates the warped segmentation mask V (i.e., segmentation mask in the source image space) by conducting inverse transformation on the target segmentation mask B ; 3) *Segmentation Module* takes the raw source image S as inputs, and output the brain segmentation mask R . The final output of our network is the extracted brain image E^M constituting only cerebral tissues, the warped image W^N aligning with the target image, and the brain segmentation mask R indicating anatomical regions of the source image.

Key Insight. Considering the highly correlated characteristics of the extraction, registration, and segmentation tasks, the design of our method revolves around the following four positive mechanisms among the three modules:

- The unalignable non-brain tissue revealed in the registration module guides the refinement process in the extraction module.
- The registration module benefits from accurate brain extraction generated in the extraction module.
- The anatomical structure estimated by the segmentation module can provide auxiliary information for the extraction module to

remove non-cerebral tissues and for the registration module to find regional correspondences.

- The segmentation module relies on the extraction and registration modules to generate the corresponding segmentation mask in the source image space.

Hence, we encode the above four positive mechanisms in our loss function design and train our network in an end-to-end fashion. Specifically, our loss function includes two perceptual losses: the image similarity term $\mathcal{L}_{sim}(W^N, T)$ and the segmentation loss term $\mathcal{L}_{seg}(R, V)$. The image similarity loss $\mathcal{L}_{sim}(\cdot, \cdot)$ captures the first and second positive mechanisms, which is affected by extraction and registration modules. The segmentation loss $\mathcal{L}_{seg}(\cdot, \cdot)$ embodies the third and fourth positive mechanisms, which is related to all modules. We jointly optimize the two loss terms to guide the learning of extraction, registration, and segmentation modules. Next, we introduce the details of each module and the training process.

3.1 Extraction Module

For the extraction module, we have adopted the multi-stage design paradigm, a strategy that has proven effective in previous works [15, 44, 45, 47, 56]. This design allows for a progressive refinement in the removal of non-cerebral tissues, culminating in an image with only cerebral tissues at the final stage. Following the multi-stage design of [45], the extraction module incorporates M extraction stages, and each stage j contains two key components as follows.

3.1.1 Extraction Network: f_e . The extraction network $f_e(\cdot)$ acts as an eliminator, intended to sequentially eradicate non-cerebral elements from source image S , aiming to retain only brain tissue in

the image obtained at the final stage. With each stage j , it forms an extraction mask $\mathbf{M}^j$, which is utilized to eliminate non-cerebral tissues in the preceding extracted brain image $\mathbf{E}^{j-1}$. Specifically, we employ the 3D U-Net [37] as the base network to learn $f_e(\cdot)$. When conducting inference, the output $\mathbf{M}^j$ is binarized by a Heaviside step function. To facilitate effective gradient backpropagation during the training phase, we employ a Sigmoid function with a large slope parameter to approximate the Heaviside step function. $f_e(\cdot)$ employs a shared-weight scheme, which means $f_e(\cdot)$ is repeatedly applied across all stages with the same set of parameters. The process can be formally expressed as:

$$\mathbf{M}^j = f_e(\mathbf{E}^{j-1}), \quad (4)$$

where $\mathbf{M}^j$ is the outputted brain mask of the j -th stage for $j = [1, \dots, M]$ and $\mathbf{E}^0 = \mathbf{S}$.

3.1.2 Overlay Layer: OL. The overlay layer serves to eradicate any residual non-cerebral tissues by applying the current brain mask $\mathbf{M}^j$ to the preceding extracted image $\mathbf{E}^{j-1}$. The updated extraction is $\mathbf{E}^j = \mathbf{E}^{j-1} \circ \mathbf{M}^j$, where $\circ$ denotes the operation of element-wise multiplication.

3.2 Registration Module

Similar to the extraction module introduced in Section 3.1, we execute a multi-stage paradigm [15, 44, 45, 47, 56] to tackle the registration task. Following the multi-stage settings of [45], the module is composed of N cascading stages, with each stage k holding the key components as follows.

3.2.1 Registration Network: f_r . The purpose of the registration network $f_r(\cdot, \cdot)$ is to slowly adapt the extracted brain image to better match the target image. Each stage k generates a current affine transformation $\mathbf{A}_i^k$ based solely on the prior warped image $\mathbf{W}^{k-1}$ and the target image $\mathbf{T}$. Similar to the extraction network f_e in Section 3.1.1, a 3D CNN-based encoder is used to learn $f_r(\cdot, \cdot)$ and a shared weight design is employed across stages using identical parameters. As a formal expression,

$$\mathbf{A}_i^k = f_r(\mathbf{W}^{k-1}, \mathbf{T}), \quad (5)$$

where $\mathbf{A}_i^k$ represents the output affine transformation of the k -th stage for $k = [1, \dots, N]$ and $\mathbf{W}^0 = \mathbf{E}^M$.

3.2.2 Composition Layer: COMP. For each stage k , upon determining the current affine transformation $\mathbf{A}_i^k$ through $f_r(\cdot, \cdot)$, we combine all preceding transformations: $\mathbf{A}_c^k = \mathbf{A}_i^k \cdot \mathbf{A}_c^{k-1}$, where $\cdot$ denoting the matrix multiplication operation. Thus, the combined transformation $\mathbf{A}_c^k$ can always be performed on the extracted image $\mathbf{E}^M$ to preserve the sharpness of the warped images [44]. For $k = 1$, the initial affine transformation $\mathbf{A}_c^0$ is designated as an identity matrix, indicating no displacement.

3.2.3 Spatial Transformation Layer: STL. A crucial process in image registration involves reconstructing the warped image $\mathbf{W}^k$ from the extracted brain image $\mathbf{E}^M$ using the affine transformation operator. Utilizing the combined transformation $\mathbf{A}_c^k$, we introduce a spatial transformation layer (STL) that resamples voxels from the extracted image to produce the warped image through $\mathbf{W}^k = \mathcal{T}(\mathbf{E}^M, \mathbf{A}_c^k)$.

Given the affine transformation operator depicted in Eq. (1), we hold

$$\mathbf{W}_{xyz}^k = \mathbf{E}_{x'y'z'}^M, \quad (6)$$

where $[x', y', z', 1]^\top = \mathbf{A}_c^k[x, y, z, 1]^\top$. To assure successful gradient propagation during this procedure, we adopt a differentiable transformation founded on trilinear interpolation proposed by [20].

3.3 Inverse Warping and Segmentation Module

This section introduces the two key designs to solve the segmentation task for the input source image $\mathbf{S}$: 1) *Inverse Warping* takes the target segmentation mask $\mathbf{B}$ and the composed affine transformation $\mathbf{A}_c^N$ as inputs, and generates a segmentation mask $\mathbf{V}$ in the source image space; 2) *Segmentation Network* takes the source image $\mathbf{S}$ as input, and predicts the brain segmentation mask $\mathbf{R}$ for $\mathbf{S}$.

3.3.1 Inverse Warping. In the training dataset, only the target image $\mathbf{T}$ is labeled with its segmentation mask $\mathbf{B}$. In order to generate the segmentation label for guiding the segmentation network, we warp the target image segmentation mask $\mathbf{B}$ into the source image space by applying the inverse of $\mathbf{A}_c^N$:

$$\mathbf{V}_{cxyz} = \mathbf{B}_{cx'y'z'}, \forall c \in \{0, 1, \dots, C-1\} \quad (7)$$

where $[x', y', z', 1]^\top = (\mathbf{A}_c^N)^{-1}[x, y, z, 1]^\top$, c is the index for anatomical class. Same as the STL in Section 3.2.3, we apply a differentiable transformation based on trilinear interpolation.

3.3.2 Segmentation Network: f_s . The segmentation network is tasked to predict a segmentation mask for the source image $\mathbf{S}$ that matches the synthesized segmentation mask $\mathbf{V}$. Similar to the extraction network discussed in Section 3.1.1, we employ the popular 3D U-Net as the base network to learn $f_s(\cdot)$. Formally, we have:

$$\mathbf{R} = f_s(\mathbf{S}). \quad (8)$$

3.4 End-to-End Training

We train our JERS model by minimizing the following objective function

$$\theta^*, \phi^*, \psi^* = \min_{\theta, \phi, \psi} \mathcal{L}_{sim}(\mathbf{W}^N, \mathbf{T}) + \lambda \mathcal{L}_{seg}(\mathbf{R}, \mathbf{V}) + \eta \sum_{j=1}^M \mathcal{R}(\mathbf{M}^j), \quad (9)$$

where θ, ϕ, ψ are the parameters for the extraction, registration, and segmentation networks, respectively. $\mathcal{L}_{sim}(\cdot, \cdot)$ is the image similarity loss, $\mathcal{L}_{seg}(\cdot, \cdot)$ is the segmentation loss, $\mathcal{R}$ is the extraction regularization term, and λ, η are positive trade-off weights.

The image similarity loss $\mathcal{L}_{sim}(\cdot, \cdot)$ is a loss function measuring the similarity between the final warped image $\mathbf{W}^N$ and the target image $\mathbf{T}$:

$$\mathcal{L}_{sim}(\mathbf{W}^N, \mathbf{T}) = \mathcal{L}_{sim}(\mathcal{T}(\mathbf{E}^M, \mathbf{A}_c^N), \mathbf{T}). \quad (10)$$

We employ the prevailing negative local cross-correlation loss for $\mathcal{L}_{sim}(\cdot, \cdot)$, which is resistant to variations in voxel intensity typically encountered across different scans and datasets [5, 51, 56]. The image similarity loss can guide the learning of both the extraction and registration modules.

The segmentation loss $\mathcal{L}_{seg}(\cdot, \cdot)$ is a loss function measuring the similarity between the predicted segmentation mask $\mathbf{R}$ and the

warped segmentation label V :

$$\mathcal{L}_{seg}(\mathbf{R}, \mathbf{V}) = \mathcal{L}_{seg}\left(f_s(\mathbf{S}), \mathcal{T}\left(\mathbf{B}, (\mathbf{A}_c^N)^{-1}\right)\right) \quad (11)$$

Here we use the cross-entropy loss function. The segmentation loss serves multiple purposes and can guide the learning of the extraction, registration, and segmentation modules.

Finally, to suppress the occurrence of other connected regions apart from the brain in the extraction masks, we define the extraction regularization loss:

$$\mathcal{R}(\mathbf{M}^j) = \sum_{x=1}^W \sum_{y=1}^H \sum_{z=1}^D \|\nabla \mathbf{M}_{xyz}^j\|^2. \quad (12)$$

This regularization term quantifies the edge strength of the predicted extraction mask $\mathbf{M}^j$, i.e., the likelihood of a voxel being an edge voxel. Through the minimization of this term, we can suppress the occurrence of edges, subsequently producing a smooth extraction mask. In this scenario, we apply the ℓ_2 -norm of the first-order derivative of $\mathbf{M}^j$ as the regularization term.

By leveraging the differentiability in each component of this design, our model can be jointly and progressively optimized in an end-to-end fashion. This training approach empowers us to unearth a joint optimal solution for the aggregate tasks of brain extraction, registration, and segmentation. Thus our work stood in marked contrast to other works [45, 51] that resort to alternative training of individual modules and can only reach a sub-optimal solution.

4 EXPERIMENTS

4.1 Datasets

The effectiveness of our proposed method is evaluated across three public real-world 3D brain MRI datasets: 1) *LPBA40* [40] includes 40 raw T1-weighted 3D MRI scans, coupled with brain masks and 56 anatomical structures as segmentation ground truths; 2) *CC359* [43] contains 359 raw T1-weighted 3D brain MRI scans and the brain masks. It also includes labeled white matter as the segmentation ground truth; 3) *IBSR* [36] consists of 18 scans with manual segmentation labels. But due to its small size, it serves only for the evaluation of the model trained on CC359. Brain masks and anatomical labels facilitate the accuracy evaluation for extraction and segmentation, respectively. Additional details can be found in Appendix A.2.

4.2 Compared Methods

We outline the comparison of our JERS with several representative methodologies in brain extraction, registration and segmentation, as illustrated in Table 1. Notably, there are no existing solutions that can seamlessly learn brain extraction, registration, and segmentation in an end-to-end framework. Hence, we designed three-stage pipeline for comparison, combining various extraction, registration, and segmentation methods. Baselines detailed in Appendix A.4.

- *Brain Extraction Tool (BET)* [41]: This technique within the FSL package uses deformable models for skull stripping.
- *3dSkullStrip (3dSS)* [10]: This is a BET variant method within the AFNI package. It uses a spherical surface to perform skull stripping.
- *Brain Surface Extractor (BSE)* [39]: This method uses morphological operations and edge detection for brain extraction, which leverages anisotropic diffusion filtering and a Marr Hildreth edge detector to identify the brain boundary.

Table 1: Summary of compared methods.

Methods	Extraction	Registration	Segmentation	Deep learning
BET [41]	✓	✗	✗	✗
3dSkullStrip [10]	✓	✗	✗	✗
BSE [39]	✓	✗	✗	✗
FLIRT [21]	✗	✓	✗	✗
ANTs [3]	✗	✓	✗	✗
VM [5]	✗	✓	✗	✓
CRN [56]	✗	✓	✗	✓
DW [20]	✗	✗	✓	✗
DeepAtlas [51]	✗	✓	✓	✓
ERNet [45]	✓	✓	✗	✓
JERS (ours)	✓	✓	✓	✓

- *FMRIB's Linear Image Registration Tool (FLIRT)* [21]: This is an automatic affine brain image registration tool in the FSL package.
- *Advanced Normalization Tools (ANTs)* [3]: This is a state-of-the-art medical image registration toolbox. We set the registration type and optimization as affine and cross-correlation metrics.
- *VoxelMorph (VM)* [5]: This is an unsupervised image registration method that uses a neural network to predict the transformation.
- *Cascaded Registration Networks (CRN)* [56]: This is an unsupervised multi-stage image registration method, which incrementally aligns the source image to the target image.
- *Directly Warping (DW)* [20]: It is an operation used to generate the segmentation mask by registration. The segmentation mask of the target image can be directly warped to the source image space by the completed registration.
- *DeepAtlas* [51]: This is a joint learning network for image registration and segmentation tasks.
- *ERNet* [45]: This is an unsupervised learning method for joint extraction and registration.

4.3 Experimental Results

Our JERS is compared with baseline methods on extraction, registration and registration accuracy. We measure the performance of each task with its corresponding metrics, and also record the time taken for each baseline method. Based on the experimental results, we find that JERS not only consistently outperforms other alternatives in terms of extraction, registration, and segmentation, but also exhibits superior time efficiency and robustness.

4.3.1 Evaluation Metrics. To evaluate the extraction and segmentation accuracy, we measure the volume overlap between the predicted and ground-truth masks. To evaluate the registration performance, we calculate the mutual information between the warped image and the target image. Details can be found in Appendix A.1.

4.3.2 Experiment Setting. We divide the datasets into training, validation and test sets. The training set is utilized for parameter learning of the model, while the validation set is employed to evaluate the performance of hyperparameter settings (e.g., the weight of the segmentation loss term). The test set is used only once to report the final evaluation results for each model. It should be noted that the IBSR dataset is exclusively used for testing purposes. We describe the detail of the data processing, JERS settings and baseline settings in Appendix A.2, A.3 and A.4. The source code is available at <https://github.com/Anonymous4545/JERS>.

Table 2: Results for brain extraction (Ext), registration (Reg) and segmentation (Seg) in different datasets. The results are reported as performance(mean $\pm$ std) of extraction, registration and segmentation of each compared method. “ $\uparrow$ ” point out “the larger the better”. The best results are highlighted in bold.

Methods			Datasets								
Ext	Reg	Seg	LPBA40			CC359			IBSR		
			Ext Dice _{ext} $\uparrow$	Reg MI $\uparrow$	Seg Dice _{seg} $\uparrow$	Ext Dice _{ext} $\uparrow$	Reg MI $\uparrow$	Seg Dice _{seg} $\uparrow$	Ext Dice _{ext} $\uparrow$	Reg MI $\uparrow$	Seg Dice _{seg} $\uparrow$
BET [41]	FLIRT [21]	DW [20]	0.935 $\pm$ 0.028	0.627 $\pm$ 0.010	0.613 $\pm$ 0.025	0.811 $\pm$ 0.087	0.481 $\pm$ 0.024	0.748 $\pm$ 0.062	0.911 $\pm$ 0.038	0.521 $\pm$ 0.022	0.800 $\pm$ 0.010
3dSS [10]	FLIRT [21]	DW [20]	0.902 $\pm$ 0.032	0.627 $\pm$ 0.010	0.601 $\pm$ 0.017	0.849 $\pm$ 0.037	0.500 $\pm$ 0.014	0.791 $\pm$ 0.034	0.869 $\pm$ 0.039	0.508 $\pm$ 0.023	0.788 $\pm$ 0.021
BSE [39]	FLIRT [21]	DW [20]	0.938 $\pm$ 0.022	0.668 $\pm$ 0.010	0.620 $\pm$ 0.008	0.846 $\pm$ 0.112	0.518 $\pm$ 0.035	0.804 $\pm$ 0.019	0.873 $\pm$ 0.064	0.521 $\pm$ 0.026	0.799 $\pm$ 0.015
BET [41]	ANTs [3]	DW [20]	0.935 $\pm$ 0.028	0.630 $\pm$ 0.013	0.625 $\pm$ 0.010	0.811 $\pm$ 0.087	0.476 $\pm$ 0.027	0.744 $\pm$ 0.062	0.911 $\pm$ 0.038	0.524 $\pm$ 0.022	0.792 $\pm$ 0.018
3dSS [10]	ANTs [3]	DW [20]	0.902 $\pm$ 0.032	0.632 $\pm$ 0.011	0.524 $\pm$ 0.133	0.849 $\pm$ 0.037	0.498 $\pm$ 0.016	0.758 $\pm$ 0.041	0.869 $\pm$ 0.039	0.508 $\pm$ 0.026	0.762 $\pm$ 0.030
BSE [39]	ANTs [3]	DW [20]	0.938 $\pm$ 0.022	0.671 $\pm$ 0.008	0.614 $\pm$ 0.023	0.846 $\pm$ 0.112	0.524 $\pm$ 0.035	0.807 $\pm$ 0.012	0.873 $\pm$ 0.064	0.522 $\pm$ 0.027	0.787 $\pm$ 0.016
BET [41]	VM [5]	DW [20]	0.935 $\pm$ 0.028	0.627 $\pm$ 0.016	0.607 $\pm$ 0.019	0.811 $\pm$ 0.087	0.465 $\pm$ 0.034	0.804 $\pm$ 0.015	0.911 $\pm$ 0.038	0.525 $\pm$ 0.023	0.795 $\pm$ 0.013
3dSS [10]	VM [5]	DW [20]	0.902 $\pm$ 0.032	0.623 $\pm$ 0.009	0.590 $\pm$ 0.009	0.849 $\pm$ 0.037	0.496 $\pm$ 0.018	0.814 $\pm$ 0.007	0.869 $\pm$ 0.039	0.509 $\pm$ 0.027	0.789 $\pm$ 0.017
BSE [39]	VM [5]	DW [20]	0.938 $\pm$ 0.022	0.670 $\pm$ 0.003	0.616 $\pm$ 0.010	0.846 $\pm$ 0.112	0.515 $\pm$ 0.041	0.805 $\pm$ 0.016	0.873 $\pm$ 0.064	0.524 $\pm$ 0.028	0.795 $\pm$ 0.013
BET [41]	CRN [56]	DW [20]	0.935 $\pm$ 0.028	0.633 $\pm$ 0.017	0.618 $\pm$ 0.022	0.811 $\pm$ 0.087	0.467 $\pm$ 0.034	0.806 $\pm$ 0.015	0.911 $\pm$ 0.038	0.527 $\pm$ 0.023	0.800 $\pm$ 0.011
3dSS [10]	CRN [56]	DW [20]	0.902 $\pm$ 0.032	0.630 $\pm$ 0.012	0.610 $\pm$ 0.013	0.849 $\pm$ 0.037	0.498 $\pm$ 0.017	0.817 $\pm$ 0.007	0.869 $\pm$ 0.039	0.513 $\pm$ 0.028	0.794 $\pm$ 0.014
BSE [39]	CRN [56]	DW [20]	0.938 $\pm$ 0.022	0.674 $\pm$ 0.006	0.626 $\pm$ 0.010	0.846 $\pm$ 0.112	0.518 $\pm$ 0.040	0.809 $\pm$ 0.017	0.873 $\pm$ 0.064	0.527 $\pm$ 0.028	0.796 $\pm$ 0.013
BET [41]	DeepAtlas [51]		0.935 $\pm$ 0.028	0.627 $\pm$ 0.016	0.645 $\pm$ 0.009	0.811 $\pm$ 0.087	0.467 $\pm$ 0.033	0.814 $\pm$ 0.017	0.911 $\pm$ 0.038	0.524 $\pm$ 0.024	0.811 $\pm$ 0.011
3dSS [10]	DeepAtlas [51]		0.902 $\pm$ 0.032	0.625 $\pm$ 0.013	0.640 $\pm$ 0.010	0.849 $\pm$ 0.037	0.498 $\pm$ 0.018	0.828 $\pm$ 0.007	0.869 $\pm$ 0.039	0.509 $\pm$ 0.027	0.808 $\pm$ 0.013
BSE [39]	DeepAtlas [51]		0.938 $\pm$ 0.022	0.667 $\pm$ 0.006	0.648 $\pm$ 0.009	0.846 $\pm$ 0.112	0.518 $\pm$ 0.041	0.817 $\pm$ 0.017	0.873 $\pm$ 0.064	0.525 $\pm$ 0.029	0.807 $\pm$ 0.014
ERNet [45]	DW [20]		0.942 $\pm$ 0.010	0.675 $\pm$ 0.004	0.620 $\pm$ 0.012	0.935 $\pm$ 0.006	0.573 $\pm$ 0.007	0.818 $\pm$ 0.007	0.915 $\pm$ 0.019	0.544 $\pm$ 0.029	0.799 $\pm$ 0.012
JERS (ours)			0.944 $\pm$ 0.008	0.679 $\pm$ 0.005	0.651 $\pm$ 0.011	0.937 $\pm$ 0.005	0.574 $\pm$ 0.006	0.840 $\pm$ 0.010	0.917 $\pm$ 0.019	0.550 $\pm$ 0.025	0.832 $\pm$ 0.010

4.3.3 Extraction, Registration and Segmentation Results. Table 2 presents the results for the compared methods as well as the proposed JERS in extraction, registration, and segmentation tasks. Through a comprehensive evaluation across three datasets, JERS outperforms existing methods in all metrics.

For the extraction task, we observed that the joint-based extraction methods (JERS and ERNet) outperform other single-stage extraction methods, especially on the CC359 dataset. Specifically, we observed a gain in extraction dice score up to 10.4% compared to the best single-stage extraction method 3dSkullStrip. Furthermore, joint-based extraction methods prove to be more robust compared to other alternatives, given their steady performance and achievement of the lowest standard deviation across all datasets.

When observing registration performance, joint-based registration methods (JERS and ERNet) also outperform all other methods across all datasets. Significantly, our findings suggest that the registration performance of most methods is constrained by the result of its corresponding extraction method. This underscores the fact that the precision of extraction considerably influences the subsequent registration task’s quality. Joint-based registration methods leverage this characteristic to yield enhanced results through collaborative learning.

For the segmentation task, once again, we find that joint-based segmentation methods (JERS and DeepAtlas) are superior to DW. This demonstrates that joint learning with the registration task can help the segmentation task to boost its performance.

Overall, joint-based methods (JERS, ERNet and DeepAtlas) outperform other pipeline-based methods in their respective joint tasks. However, the partially joint methods perform poorly on their stand-alone task. ERNet performs well on extraction and registration but inferior on the segmentation task. Similarly, although DeepAtlas achieves good results on the segmentation task, the stand-alone extraction method limits its extraction and registration performance. Benefiting from fully end-to-end joint learning, only JERS can perform well in all tasks.

4.3.4 Qualitative Analysis. In Figure 5, we visually compare the performance of our JERS and other approaches on the LPBA test set. Upon observation, it is evident that JERS achieves more accurate brain extraction compared to BET, 3dSkullStrip, and BSE. The brain extraction mask predicted by JERS overlaps closely with the ground truth extraction mask, while the masks predicted by other extraction methods include noticeable non-brain tissues. Regarding registration results, JERS also outperforms the other methods. The final registered image of JERS exhibits a higher resemblance to the target image compared to the alternatives. Importantly, the inaccurate extraction results with non-brain tissue further impact the subsequent registration results and ultimately affect the overall performance. This supports our assertion that failed extraction can propagate errors to the subsequent registration task, resulting in irreparable consequences. In terms of the segmentation task, JERS produces results that closely overlap to the ground truth segmentation mask. Overall, we can see that better extraction leads to better registration, and better registration yields better segmentation.

4.3.5 Ablation Study. To demonstrate the effectiveness of our JERS, we compared five variants of JERS in Table 3. We first freeze the extraction module, the registration module, and the segmentation module of JERS respectively. JERS w/o Ext consistently produces the extraction mask with all values of 1 (*i.e.*, no pattern be removed from the source image). JERS w/o Reg only outputs the identity affine matrix representing no displacement applied to the image. JERS w/o Seg directly warps the segmentation mask to the source image space (*i.e.*, no segmentation network exists). The results show that the extraction module, the registration module, and the segmentation module are essential to JERS, and removing any of them degrades the performance of all tasks. We then evaluate the effects of multi-stage and extraction mask smoothing design. The results show that they significantly boost the performance of all tasks. Since all tasks are learned in a collective manner, the performance boost of one module is shared by all other modules.

Methods		BET - FLIRT - DW		3dSS - ANTs - DW		BSE - VM - DW		BET - CRN - DW		3dSS - DeepAtlas		ERNet - DW		JERS (ours)	
Input & Evaluation Data															
Template	Source	Ext	Reg	Ext	Reg	Ext	Reg	Ext	Reg	Ext	Reg	Ext	Reg	Ext	Reg
															
															
															
															
Template Mask	Source	Ext Mask	Seg Mask	Ext Mask	Seg Mask	Ext Mask	Seg Mask	Ext Mask	Seg Mask	Ext Mask	Seg Mask	Ext Mask	Seg Mask	Ext Mask	Seg Mask
															
															
															
															

Figure 5: Visual comparisons for brain extraction, registration and segmentation tasks. We render a 3D visualization of the image and display the middle slice in three different planes: sagittal, axial and coronal. The left side contains the source and target (template) images and their corresponding ground truth labels. We show the extraction, registration and segmentation results of each method and its corresponding predictive labels used for performance evaluations. For the extraction task, a predicted extraction mask (marked by green color) should coincide as much as possible with the ground truth extraction mask of the source image. Likewise, in the segmentation task, a predicted segmentation mask (marked by different color regions) should well-overlap with the ground-truth segmentation mask of the source image. For the registration task, the higher the similarity of the registered brain to the template brain, the better.

4.3.6 Running Efficiency. We measure the efficiency of JERS by comparing its running time with other baselines. The measurement is made on the same device with an Intel® Xeon® E5-2667 v4 CPU and an NVIDIA Tesla A100 GPU. As indicated in Table 4, all joint-based methods are faster than existing three-stage pipeline-based methods. This is because they can efficiently perform their corresponding extraction, registration and segmentation tasks end-to-end on the same device. GPU implementations for BET, 3dSkullStrip, BSE, FLIRT, and ANTs are not available [3, 10, 21, 39, 41].

4.3.7 Influence of Parameters. We study two crucial hyperparameters of our JERS: the number of stages for extraction and registration modules and the value of segmentation loss weight λ . In our multi-stage design, the number of stages in the network represents the depth of the model and the number of iterations for the extraction and registration tasks. Essentially, increasing the number of stages allows for more refinements in the extraction and registration processes. As illustrated in Figure 6(a, b, c), we adjust the number of stages of extraction and registration to study their influence. The outcomes show that increasing the stages boosts the performance of all tasks, validating the notion that a multi-stage framework results in superior overall performance in a joint learning system.

As mentioned in Section 3.4, we introduce a segmentation loss term to learn a better segmentation network. To show the effectiveness of the loss term, we vary different values of the loss weight λ as shown in Figure 6(d). As the weight of the loss term gradually increases, the segmentation dice score grows as well. This indicates that our JERS benefits from the segmentation loss term.

5 RELATED WORK

Neuroimage extraction. In the past decade, numerous methods have emerged, highlighting the significance of the brain extraction problem. Smith et al. [41] introduced a deformable model that fits the brain surface using a locally adaptive set model. 3dSkullStrip [10] is a modified version of [41] that employs points outside the brain surface to guide mesh evolution. Shattuck et al. [39] utilized anisotropic diffusion filtering and a 2D Marr Hildreth edge detector for brain boundary identification. However, these methods heavily rely on parameter tuning and manual quality control, which are time-consuming and labor-intensive. Recently, brain extraction has benefited from the introduction of deep learning approaches, which exhibit exceptional performance and speed.

Table 3: Ablation studies of JERS on LPBA40 dataset


Methods	Dice _{ext} ↑	MI _{reg} ↑	Dice _{seg} ↑
JERS w/o Ext	0.216 ± 0.018	0.579 ± 0.017	0.224 ± 0.103
JERS w/o Reg	0.311 ± 0.030	0.217 ± 0.036	0.268 ± 0.157
JERS w/o Seg	0.942 ± 0.009	0.677 ± 0.005	0.620 ± 0.008
JERS w/o Multi-stage	0.902 ± 0.005	0.662 ± 0.003	0.609 ± 0.017
JERS w/o Ext smoothing	0.931 ± 0.006	0.673 ± 0.003	0.649 ± 0.008
JERS	0.944 ± 0.008	0.679 ± 0.005	0.651 ± 0.011

Table 4: Running Time of compared methods on LPBA40 dataset.

Methods			Time (Sec) ↓		
Ext	Reg	Seg	Ext	Reg	Seg
BET [41]	FLIRT [21]	DW [20]	2.4751	4.6857	0.1457
3dSS [10]	ANTs [3]	DW [20]	176.9355	2.6705	0.0370
BSE [39]	VM [5]	DW [20]	3.5119	0.0050	0.0001
BET [41]	CRN [56]	DW [20]	2.4751	0.0151	0.0001
BET [41]	DeepAtlas [51]		2.4751	0.0018	
ERNet [45]	DW [20]		0.0420	0.0001	
JERS (ours)			0.0455		

Kleesiek et al. [23] proposed a voxel-wise 3D CNN for skull stripping, while Hwang et al. [19] demonstrated the effectiveness of 3D-UNet in achieving competitive results. However, these learning-based approaches often necessitate a substantial amount of properly labeled data for effective training, which is a challenge considering that neuroimage datasets are typically small and costly to annotate. **Neuroimage registration.** Conventional techniques for image registration [2, 3, 21] typically aim to maximize image similarity by iteratively optimizing transformation parameters. Commonly used intensity-based similarity measures include normalized cross-correlation (NCC) and mutual information (MI), among others. However, this iterative optimization approach often suffers from high computational costs and being stuck in local optima, resulting in inefficient and unreliable registration outcomes. Recently, numerous deep learning-based methods have been proposed, offering improved computational efficiency and registration performance. For instance, Sokooti et al. [42] introduced a multi-scale 3D CNN called RegNet, which learns the displacement vector field (DVF) for 3D chest CT registration. Although these methods demonstrate competitive results, they require supervision. To overcome this limitation, unsupervised registration methods [5, 56] have garnered significant attention and shown promising outcomes.

Neuroimage segmentation. CNN-based approaches have demonstrated superior performance in terms of speed and accuracy for supervised neuroimage segmentation. Based on the dimensionality of network operation, current work mainly falls into two categories: 2D CNN-based network and 3D CNN-based network. 2D CNN-based network processes the volumetric neuroimage data slice by slice, and assembles the final segmentation mask by putting segmentation results on all the slices together. The most representative networks include UNet [37] and Attention-UNet [32]. Since the 2D CNN-based network treats each slice separately, the spatial information encoded between slices is not utilized, negatively affecting the segmentation performance. This limitation prompts the use of the 3D CNN-based network for 3D segmentation. The most representative network among all is the 3D UNet [9]. Despite the success CNN-based approaches had in neuroimage segmentation, those approaches still possess a data-hungry nature and require a large number of labeled neuroimages for training, which is expensive to acquire.

**Figure 6: Effect of varying the number of stages of the JERS and the segmentation loss weight λ .**

Joint neuroimage registration and segmentation The neuroimage registration and segmentation tasks are deeply co-related and mutually facilitating, thus should not be treated separately. Specifically, the anatomy structure produced by neuroimage segmentation can provide auxiliary information for neuroimage segmentation. This idea is practiced in [17] and demonstrated to be effective. Complementary, the registration task can also aid the segmentation task, typically studied in the scope of atlas-based segmentation and data augmentation. In atlas-based segmentation, the atlas label is transferred to the unlabeled image space using the geometric transformation estimated by a registration module [14, 49]. Taking the idea a step further, for data augmentation, new image-label pairs are generated by means of sampling the geometric and style transformation [13, 55]. The above methods still focus on a single task. To address this limitation, exploring work has been conducted to combine registration and segmentation together, by optimizing them either jointly [35] or alternatively [16, 51], and is able to obtain better results.

6 CONCLUSION

This paper introduces a novel unified framework, called JERS, for one-shot joint extraction, registration and segmentation. In contrast to prior research, our proposed method seamlessly integrates the three tasks into a single system, enabling joint optimization. Specifically, JERS contains a group of collective extraction, registration and segmentation modules. These three modules help each other boost extraction, registration, and segmentation performance with only one labeled template available. Furthermore, our method facilitates incremental progress for each task, thereby enhancing the overall performance to a greater extent. The experimental results substantiate that JERS not only surpasses state-of-the-art approaches in terms of extraction, registration, and segmentation accuracy but also exhibits high robustness and time-efficiency.

7 ACKNOWLEDGMENTS

This work is partially supported by the National Science Foundation (MRI 2215789 and IIS 1909879) and Lehigh's grants (S00010293 and 001250).

REFERENCES

- [1] Zeynettin Akkus, Alfia Galimzianova, Assaf Hoogi, Daniel L Rubin, and Bradley J Erickson. 2017. Deep learning for brain MRI segmentation: state of the art and future directions. *Journal of digital imaging* (2017).
- [2] Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *MIA* (2008).
- [3] Brian B Avants, Nick Tustison, and Gang Song. 2009. Advanced normalization tools (ANTS). *Insight j* (2009).
- [4] Zilong Bai, Peter Walker, Anna Tschiffely, Fei Wang, and Ian Davidson. 2017. Unsupervised network discovery for brain imaging data. In *SIGKDD*.
- [5] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. 2018. An unsupervised learning model for deformable medical image registration. In *CVPR*.
- [6] Lei Cai, Zhengyang Wang, Hongyang Gao, Dinggang Shen, and Shuiwang Ji. 2018. Deep adversarial learning for multi-modality missing data completion. In *SIGKDD*.
- [7] Xu Chen, Bryan M Williams, Srinivasa R Vallabhaneni, Gabriela Czanner, Rachel Williams, and Yalin Zheng. 2019. Learning active contour models for medical image segmentation. In *CVPR*.
- [8] Yongjun Chen, Hongyang Gao, Lei Cai, Min Shi, Dinggang Shen, and Shuiwang Ji. 2018. Voxel deconvolutional networks for 3D brain image labeling. In *SIGKDD*.
- [9] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 2016. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*.
- [10] Robert W Cox. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research* (1996).
- [11] Xin Dai, Xiangnan Kong, Tian Guo, John Boaz Lee, Xinyue Liu, and Constance Moore. 2020. Recurrent networks for guided multi-attention classification. In *SIGKDD*.
- [12] Xin Dai, Xiangnan Kong, Xinyue Liu, John Boaz Lee, and Constance Moore. 2020. Dual-Attention Recurrent Networks for Affine Registration of Neuroimaging Data. In *SDM*.
- [13] Yuhang Ding, Xin Yu, and Yi Yang. 2021. Modeling the probabilistic distribution of unlabeled data for one-shot medical image segmentation. In *AAAI*.
- [14] Nicola K Dinsdale, Mark Jenkinson, and Ana IL Namburete. 2019. Spatial warping network for 3D segmentation of the hippocampus in MR images. In *MICCAI*.
- [15] Wanquan Feng, Juyong Zhang, Hongrui Cai, Haoqi Xu, Junhui Hou, and Huijun Bao. 2021. Recurrent multi-view alignment network for unsupervised surface registration. In *CVPR*.
- [16] Yuting He, Tiantian Li, Guanyu Yang, Youyong Kong, Yang Chen, Huazhong Shu, Jean-Louis Coatrieux, Jean-Louis Dillenseger, and Shuo Li. 2020. Deep complementary joint model for complex scene registration and few-shot segmentation on medical images. In *ECCV*.
- [17] Yipeng Hu, Marc Modat, Eli Gibson, Wenqi Li, Nooshin Ghavami, Ester Bonmati, Guotai Wang, Steven Bandula, Caroline M Moore, Mark Emberton, et al. 2018. Weakly-supervised convolutional neural networks for multimodal image registration. *MIA* (2018).
- [18] Shuai Huang, Jing Li, Jieping Ye, Adam Fleisher, Kewei Chen, Teresa Wu, and Eric Reiman. 2011. Brain effective connectivity modeling for Alzheimer's disease by sparse Gaussian Bayesian network. In *SIGKDD*.
- [19] Hyunho Hwang, Hafiz Zia Ur Rehman, and Sungon Lee. 2019. 3D U-Net for skull stripping in brain MRI. *Applied Sciences* (2019).
- [20] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and koray kavukcuoglu. 2015. Spatial Transformer Networks. In *NeurIPS*.
- [21] Mark Jenkinson and Stephen Smith. 2001. A global optimisation method for robust affine registration of brain images. *MIA* (2001).
- [22] Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *MIA* (2017).
- [23] Jens Kleesiek, Gregor Urban, Alexander Hubert, Daniel Schwarz, Klaus Maier-Hein, Martin Bendszus, and Armin Biller. 2016. Deep MRI brain extraction: A 3D convolutional neural network for skull stripping. *NeuroImage* (2016).
- [24] John Boaz Lee, Xiangnan Kong, Yihan Bao, and Constance Moore. 2017. Identifying deep contrasting networks from time series data: Application to brain network analysis. In *SDM*.
- [25] John Boaz Lee, Xiangnan Kong, Constance M Moore, and Nesreen K Ahmed. 2020. Deep Parametric Model for Discovering Group-cohesive Functional Brain Regions. In *SDM*.
- [26] Sikun Lin, Shuyun Tang, Scott T. Grafton, and Ambuj K. Singh. 2022. Deep Representations for Time-Varying Brain Datasets. In *SIGKDD*.
- [27] Xinyue Liu, Xiangnan Kong, and S Yu Philip. 2017. Collective discovery of brain networks with unknown groups. In *IJCNN*.
- [28] Xinyue Liu, Xiangnan Kong, and Ann B Ragin. 2017. Unified and contrasting graphical lasso for brain network discovery. In *SDM*.
- [29] Oeslle Lucena, Roberto Souza, Leticia Rittner, Richard Frayne, and Roberto Lotufo. 2019. Convolutional neural networks for skull-stripping in brain MR imaging using silver standard masks. *AIM* (2019).
- [30] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. 1997. Multimodality image registration by maximization of mutual information. *TMI* (1997).
- [31] Frederik Maes, Dirk Vandermeulen, and Paul Suetens. 2003. Medical image registration using mutual information. *Proc. IEEE* (2003).
- [32] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. 2018. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999* (2018).
- [33] Evangelos E Papalexakis, Alona Fyshe, Nicholas D Sidiropoulos, Partha Pratim Talukdar, Tom M Mitchell, and Christos Faloutsos. 2014. Good-enough brain model: Challenges, algorithms and discoveries in multi-subject experiments. In *SIGKDD*.
- [34] Josien PW Pluim, JB Antoine Maintz, and Max A Viergever. 2000. Image registration by maximization of combined mutual information and gradient information. In *MICCAI*.
- [35] Liang Qiu and Hongliang Ren. 2021. U-RSNet: An unsupervised probabilistic model for joint registration and segmentation. *Neurocomputing* (2021).
- [36] Torsten Rohlfing. 2011. Image similarity and tissue overlaps as surrogates for image registration accuracy: widely used but unreliable. *TMI* (2011).
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*.
- [38] Florent Ségonne, Anders M Dale, Evelina Busa, Maureen Glessner, David Salat, Horst K Hahn, and Bruce Fischl. 2004. A hybrid approach to the skull stripping problem in MRI. *Neuroimage* (2004).
- [39] David W Shattuck and Richard M Leahy. 2002. BrainSuite: an automated cortical surface identification tool. *MIA* (2002).
- [40] David W Shattuck, Mubeena Mirza, Vitria Adisetiyo, Cornelius Hojatkashani, Georges Salamon, Katherine L Narr, Russell A Poldrack, Robert M Bilder, and Arthur W Toga. 2008. Construction of a 3D probabilistic atlas of human cortical structures. *Neuroimage* (2008).
- [41] Stephen M Smith. 2002. Fast robust automated brain extraction. *Human brain mapping* (2002).
- [42] Hessam Sokooti, Bob De Vos, Floris Berendsen, Boudewijn PF Lelieveldt, Ivana Išgum, and Marius Staring. 2017. Nonrigid image registration using multi-scale 3D convolutional neural networks. In *MICCAI*.
- [43] Roberto Souza, Oeslle Lucena, Julia Garrafa, David Gobbi, Marina Saluzzi, Simone Appenzeller, Leticia Rittner, Richard Frayne, and Roberto Lotufo. 2018. An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement. (2018).
- [44] Yao Su, , Xin Dai, Lifang He, and Xiangnan Kong. 2022. ABN: Anti-Blur Neural Networks for Multi-Stage Deformable Image Registration. In *ICDM*.
- [45] Yao Su, Zhentian Qian, Lifang He, and Xiangnan Kong. 2022. Ernet: Unsupervised collective extraction and registration in neuroimaging data. In *SIGKDD*.
- [46] Liang Sun, Rinkal Patel, Jun Liu, Kewei Chen, Teresa Wu, Jing Li, Eric Reiman, and Jieping Ye. 2009. Mining brain region connectivity for alzheimer's disease study via sparse inverse covariance estimation. In *SIGKDD*.
- [47] Yujiao Tang, Feng Yang, Shaofeng Yuan, et al. 2019. A multi-stage framework with context information fusion structure for skin lesion segmentation. In *ISBI*.
- [48] Philippe Thévenaz and Michael Unser. 2000. Optimization of mutual information for multiresolution image registration. *TIP* (2000).
- [49] Shuxin Wang, Shilei Cao, Dong Wei, Renzhen Wang, Kai Ma, Liansheng Wang, Deyu Meng, and Yefeng Zheng. 2020. LT-Net: Label transfer by learning reversible voxel-wise correspondence for one-shot medical image segmentation. In *CVPR*.
- [50] Shen Wang, Lifang He, Bokai Cao, Chun-Ta Lu, Philip S Yu, and Ann B Ragin. 2017. Structural deep brain network mining. In *SIGKDD*.
- [51] Zhenlin Xu and Marc Niethammer. 2019. DeepAtlas: Joint semi-supervised learning of image registration and segmentation. In *MICCAI*.
- [52] Sen Yang, Qian Sun, Shuiwang Ji, Peter Wonka, Ian Davidson, and Jieping Ye. 2015. Structural graphical lasso for learning mouse brain connectivity. In *SIGKDD*.
- [53] Hang Yin, Xiangnan Kong, and Xinyue Liu. 2018. Coherent graphical lasso for brain network discovery. In *ICDM*.
- [54] Hang Yin, Xinyue Liu, and Xiangnan Kong. 2020. Gaussian mixture graphical lasso with application to edge detection in brain networks. In *BigData*.
- [55] Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. 2019. Data augmentation using learned transformations for one-shot medical image segmentation. In *CVPR*.
- [56] Shengyu Zhao, Yue Dong, Eric I Chang, Yan Xu, et al. 2019. Recursive cascaded networks for unsupervised medical image registration. In *ICCV*.

A APPENDIX

This section provides more details of evaluation metrics and experiment settings to support the reproducibility of the results in this paper. Our code and data have been made publicly available at <https://github.com/Anonymous4545/JERS>.

A.1 Evaluation Metrics

Our defined problem aims to solve the brain extraction, registration and segmentation tasks simultaneously: 1) identify the brain region (*i.e.*, whole cerebral tissue) within the source image; 2) align the extracted cerebral tissues to the target image; 3) identify the anatomical segmentation within the source image. Thus, we evaluate the accuracy of extraction, registration and segmentation to show the performance of our proposed method and compared methods.

A.1.1 Extraction Performance. The brain MRI datasets incorporate the ground truth of the brain mask, representing the labeling of brain tissue in the source image. To assess the accuracy of extraction, we evaluate the volume overlap of brain masks using the Dice score, which can be expressed as:

$$\text{Dice}_{\text{ext}} = 2 \cdot \frac{|\hat{\mathbf{M}} \cap \mathbf{M}|}{|\hat{\mathbf{M}}| + |\mathbf{M}|}, \quad (13)$$

where $\hat{\mathbf{M}}$ denotes the predicted brain mask and $\mathbf{M}$ is the corresponding ground truth. If $\hat{\mathbf{M}}$ signifies a precise extraction, we anticipate a high degree of overlap between the non-zero areas in $\hat{\mathbf{M}}$ and $\mathbf{M}$.

A.1.2 Registration Performance. To evaluate the registration performance, we calculate the mutual information [30, 31, 34, 48] between the warped image (*i.e.*, registered) $\mathbf{W}$ and the target image $\mathbf{T}$:

$$MI(\mathbf{W}, \mathbf{T}) = \sum_{w,t} p_{\mathbf{WT}}(w, t) \log \frac{p_{\mathbf{WT}}(w, t)}{p_{\mathbf{W}}(w) \cdot p_{\mathbf{T}}(t)} \quad (14)$$

where $p_{\mathbf{W}}(w)$ and $p_{\mathbf{T}}(t)$ are the marginal probability distributions of image $\mathbf{W}$ and $\mathbf{T}$, respectively. $p_{\mathbf{WT}}(w, t)$ is the joint probability distribution. The mutual information measures the mutual dependence between $\mathbf{W}$ and $\mathbf{T}$. If the warped image $\mathbf{W}$ and the target image $\mathbf{T}$ are geometrically aligned, we expect the mutual information to be maximal.

A.1.3 Segmentation Performance. We evaluate the segmentation accuracy by measuring the volume overlap of anatomical segmentation, which are the location labels of different tissues in the brain MRI image. If the segmentation task performs well, the predicted segmentation mask should overlap with the ground truth. Similar to the extraction evaluation, we use Dice score to evaluate the overlap of the segmentation masks. A Dice score of 1 signifies that the corresponding structures overlap with the ground truth, whereas a score of 0 denotes the complete absence of overlap. If the image includes multiple labeled anatomical structures, the final score is calculated as the average of the Dice scores for each structure.

A.2 Details of Data Preprocessing

The proposed method and baselines are evaluated on three different public brain MRI datasets, LPBA40, CC-359 and IBSR.

- *LONI Probabilistic Brain Atlas (LPBA40)* [40]: The dataset comprises 40 raw T1-weighted 3D brain MRI scans, each from a different patient. It includes corresponding brain masks and segmentation

ground truth for 56 anatomical structures. The brain mask is used for evaluating the extraction accuracy, while the anatomical segmentations are used for evaluating the segmentation accuracy. Similar to [5, 56], our focus is on atlas-based registration, where the first scan serves as the target image and the remaining scans are aligned to it. Out of the 39 remaining scans, 30 are used for training, 5 for validation, and 4 for testing. All scans are cropped and resized to $96 \times 96 \times 96$ dimensions.

- *Calgary-Campinas-359 (CC-359)* [43]: The dataset consists of 359 raw T1-weighted 3D brain MRI scans from 359 different patients. Additionally, it includes corresponding brain masks and labeled white matter as ground truth. The brain masks are used to evaluate the accuracy of extraction, while the white matter masks are used for segmentation evaluation. Similar to the LPBA40 dataset, our focus is on atlas-based registration. For CC359, we divide the dataset into training, validation, and test sets, consisting of 298, 30, and 30 scans, respectively. All scans are cropped and resized to $96 \times 96 \times 96$.

- *Internet Brain Segmentation Repository (IBSR)* [36]: The dataset consists of 18 raw T1-weighted 3D brain MRI scans from 18 different patients, accompanied by corresponding segmentation results. The segmentation results are merged to create the brain mask. Given the limited sample size, this dataset is exclusively used for testing the model trained on CC359. Consequently, all 18 scans are aligned with the first scan of CC359. After cropping, all scans are resized to $96 \times 96 \times 96$ dimensions.

A.3 Details Settings of JERS

Training settings of JERS. Our experiments are conducted on Red Hat Enterprise Linux 7.3, utilizing an Intel® Xeon® E5-2667 v4 CPU and an NVIDIA Tesla A100 GPU. The code is implemented in Python 3.7.6, and the neural networks are built using PyTorch 1.7.1. The implementation also makes use of Numpy 1.21.6, SimpleITK 2.0.2, and Nibabel 3.1.1. To overcome GPU memory limitations, we employ batch gradient descent, with each training batch consisting of one image pair. The models are optimized using the Adam optimizer, with a learning rate of 1×10^{-6} . We also apply image augmentation techniques, including random translation, rotation, and scaling, to the source images during training. For more details, please refer to Table 5.

Table 5: Range of random transformation.

Datasets	Transformation		
	Translation (Voxels)	Rotation (Degree)	Scale (Times)
LPBA40	± 5	± 5	0.98 ~ 1.02
CC359	± 3	± 3	0.99 ~ 1.01

Parameters settings of JERS. The extraction and registration stages are set to 5 in this work. The segmentation loss parameter λ and extraction mask smooth parameter η in Eq. (9) are 0.1 and 1, respectively. The extraction network contains 10 convolutional layers with 16, 32, 32, 64, 64, 64, 32, 32, 32 and 16 filters. The registration network adopt 3D CNNs and fully-connected layers to map the input to the dimension of 1×12 . It contains 6 convolutional layers with 16, 32, 64, 128, 256 and 512 filters. The segmentation network contains 10 convolutional layers with 128, 256, 256, 512, 512, 512, 256, 256, 256 and 128 filters.

A.4 Settings of Baselines

The settings of baselines are followed by [45] for a fair comparison.

Brain Extraction Tool (BET) [41]: This skull stripping method is a component of the FSL (FMRIB Software Library) package. It employs a deformable approach to accurately fit the brain surface by utilizing locally adaptive set models. The command we use for BET is `bet <input> <output> -f 0.5 -g 0 -m`, where `f` and `g` are fractional intensity threshold and gradient in fractional intensity threshold, respectively. We set them to default values.

3dSkullStrip [10]: This modified version of BET (Brain Extraction Tool) is integrated into the AFNI (Analysis of Functional NeuroImages) package. It performs skull stripping by employing the expansion paradigm of the spherical surface. The command we use for 3dSkullStrip is `3dSkullStrip -input <input> -prefix <output> -mask_vol -fac 1000`. `fac` is set to the default value.

Brain Surface Extractor (BSE) [39]: This method extracts the brain region by utilizing morphological operations and edge detection techniques. It incorporates anisotropic diffusion filtering to enhance image quality and a Marr Hildreth edge detector to accurately identify the boundaries of the brain. The command we use for BSE is `bse -i <input> -o <output> -mask <mask> -p -trim -auto -timer`. Hyperparameters are set to default values.

FMRIB's Linear Image Registration Tool (FLIRT) [21]: This is a fully automated affine brain image registration tool included in the FSL (FMRIB Software Library) package. It performs the registration process without requiring manual intervention, allowing for the alignment of brain images based on affine transformations. The command we use for FLIRT is `flirt -in <source> -ref <target> -out <output> -omat <output parameter> -bins 256 -cost corratio -searchrx -90 90 -searchry -90 90 -searchrz -90 90 -dof 12 -interp trilinear`.

Advanced Normalization Tools (ANTs) [3]: It is a cutting-edge medical image registration toolkit widely used in the field. In our

approach, we employ the affine transformation model and cross-correlation metric provided by ANTs for the registration process.

VoxelMorph (VM) [5]: This unsupervised image registration method utilizes a neural network to predict the transformation between images. In order to ensure a fair comparison, we re-implemented the method using an affine transformation. The network architecture consists of 6 convolutional layers with filter sizes of 16, 32, 64, 128, 256, and 512. The deformation regularization ratio is set to 10, ensuring smooth and controlled transformations during the registration process.

Directly Warping (DW) [20]: This operation refers to generating a segmentation mask through the process of registration. Once the registration is completed, the segmentation mask of the target image can be directly warped and transformed into the source image space.

Cascaded Registration Networks (CRN) [56]: This is an unsupervised multi-stage registration method that involves iteratively transforming the source image to align with a target image. Same to JERS, the number of stages is set to 5. Within each stage, we configure the network architecture with 6 convolutional layers using filter sizes of 16, 32, 64, 128, 256, and 512.

DeepAtlas [51]: This is an unsupervised learning method for joint registration and segmentation. For a fair comparison, we configure 6 convolutional layers with 16, 32, 64, 128, 256 and 512 filters for the registration module, and the segmentation network contains 10 convolutional layers with 128, 256, 256, 512, 512, 512, 256, 256, 256 and 128 filters.

ERNet [45]: This is an unsupervised learning method for joint extraction and registration. For a fair comparison, the number of stages is set to 5, and we configure 6 convolutional layers with 16, 32, 64, 128, 256 and 512 filters, and the registration network contains 6 convolutional layers with 16, 32, 64, 128, 256 and 512 filters.