

TemplateFlow: a community archive of imaging templates and atlases for improved consistency in neuroimaging

R Ciric¹, R Lorenz^{1,2,3}, WH Thompson⁴, M Goncalves¹, E MacNicol⁵, CJ Markiewicz¹, YO Halchenko⁶, SS Ghosh^{7,8}, KJ Gorgolewski¹, RA Poldrack¹, and O Esteban^{1,9}

¹Department of Psychology, Stanford University, CA, USA; ²MRC CBU, University of Cambridge, Cambridge, UK; ³Department of Neurophysics, MPI, Leipzig, Germany; ⁴Department of Clinical Neuroscience, Karolinska Institutet, Stockholm, Sweden; ⁵Department of Neuroimaging, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK; ⁶Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH, USA. ⁷McGovern Institute for Brain Research, Massachusetts Institute of Technology: MIT, Cambridge, MA, USA; ⁸Department of Otolaryngology, Harvard Medical School, Boston, MA, USA; ⁹Department of Radiology, University Hospital of Lausanne and University of Lausanne, Lausanne, Switzerland.

Neuroimaging templates and corresponding atlases play a central role in experimental workflows and are the foundation for reporting of results. The proliferation of templates and atlases is one relevant source of methodological variability across studies, which has been brought to attention recently as an important challenge to reproducibility in neuroscience. Unclear nomenclature, an overabundance of template variants and options, inadequate provenance tracking and maintenance, and poor concordance between atlases introduce further unreliability into reported results. We introduce **TEMPLATEFLOW**, a cloud-based repository of human and nonhuman brain templates paired with a client application for programmatically accessing resources. **TEMPLATEFLOW** is designed to be extensible, providing a transparent pathway for researchers to contribute and vet templates and their associated atlases. Following software engineering best practices, **TEMPLATEFLOW** leverages technologies for unambiguous resource identification, data management, versioning and synchronisation, programmatic extensibility, and continuous integration. By equipping researchers with a robust resource for using and evaluating brain templates, **TEMPLATEFLOW** will contribute to increasing the reliability of neuroimaging results.

Introduction

Brains are morphologically variable, exhibiting diversity in such features as overall size¹, sulcal curvature², and functional topology^{3,4}. Morphological variability manifests not only in differences between brains, but also in the way that a brain changes across its lifespan, as it is remodelled by development, ageing, and degenerative processes^{5–7}. These morphological

differences correspond with effects of interest in neuroimaging studies but also hinder direct comparisons between images of brains⁸. The substantial variability within and between individual brains necessitates a means of formalising population-level knowledge about brain anatomy and function. Neuroscientists have answered this need by creating brain atlases as references for understanding and contextualising morphological variability. Atlases encapsulate landmarks, features, and other kinds of knowledge about the brain as annotations that are consistent across individual brains.

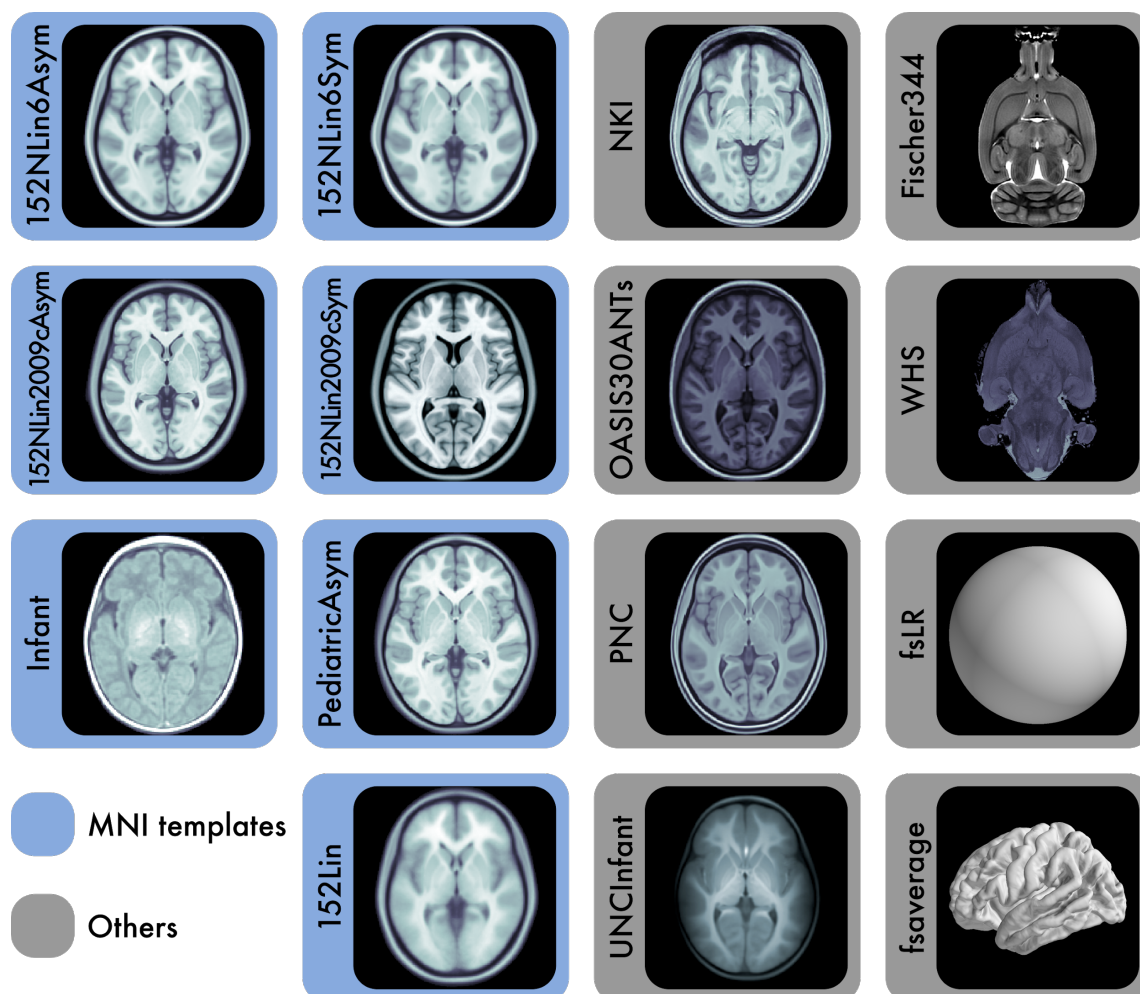
The development of atlases in neuroscience has accelerated knowledge discovery and dissemination. Early endeavours, epitomised by the groundbreaking work of Brodmann, leveraged careful scrutiny of microanatomy and cytoarchitectonic properties in small numbers of brains^{9,10}. Concurrent macroanatomical approaches, by contrast, identified common features in nuclear boundaries and cortical gyrification. Modern atlases advanced on these approaches by incorporating stereotaxy¹¹, defining a basis set of coordinate axes over the brain and anchoring neural landmarks to coordinates. Initially developed over a century ago to surgically induce targeted brain lesions, the first stereotaxic apparatus informed early sectional atlases of the cat and macaque brains¹¹. In humans, Talairach's assiduous postmortem examination of a single human brain produced a stereotaxic atlas that saw wide use¹². Since then, neuroscientists have directed great efforts to improve existing¹³ and generate new atlases of the neurotypical adult human¹⁴ and nonhuman^{15,16} brain; as well as developing, ageing and neurologically atypical brains. For instance, new atlases and representative stereotaxic maps can be created for diseased¹⁷, infant^{18,19}, and elderly²⁰ human populations or to capture the rapid postnatal development of nonhuman species^{21,22}. Recently, atlas endeavours have largely shifted away from the search for a single universal neuroanatomical pattern, instead making use of increasingly large samples with the aim of representing a population average of the distribution of morphological patterns.

On account of its relatively high spatial resolution, its capacity to image the entire brain, and its non-invasive acquisition protocols, magnetic resonance imaging (MRI) has revolutionised neuroscience in general, and the atlas endeavours²³ in particular. For three cardinal reasons, atlases have become an indispensable component of modern neuroimaging data workflows. First, group inference in neuroimaging studies requires that individuals' features are aligned into a common spatial frame of reference where their location can be called standard⁸. Second, the progress of software instruments to map homologous features between subjects^{24,25} has enabled researchers to create population-average maps of a particular image modality with relative ease using commonly available software. These maps, called *templates*, are typically created by averaging images that are representative of the population of interest to a study^{26,27}. Concomitant advances in image acquisition, processing, and analysis have enabled new templates to build iteratively upon previous work²³. Third, templates engender a stereotaxic coordinate system in which atlases can be delineated or projected. Associating atlases with template coordinates also facilitates the mapping of prior population-level knowledge about the brain into images of individual subjects' brains (for instance, to sample and average the functional MRI signal indexed by the regions defined in an atlas²⁸).

Because they are integral to analytic workflows and because atlas technology is continuously improving, a multiplicity of brain templates and atlases have been published²⁷. Factors that have stimulated the proliferation of templates include distribution within software toolboxes, data structuring conventions, maintenance, and issues with licensing and shareability. As a result, researchers have at their disposal a wealth of open-access templates and atlases as well as established protocols for the creation of study-specific alternatives. Recent research²⁹ has reappraised the naive *more-is-better* assumption for methodological options, highlighting how more options and greater methodological flexibility can threaten the reproducibility of findings. This problem is evidently not new for the neuroimaging community³⁰, and the proliferation of templates and atlases only adds on to the methodological degrees of

freedom available to researchers. In the particular case of templates and atlases, the problem is exacerbated by poor consistency across available alternatives, as exposed by Bohland et al.³¹. Along similar lines, Yoon et al. cautioned about “template effects” confounding the interpretation of results from pediatric imaging studies based on a common adult reference template²⁶. The concerning picture is completed with inaccurate reporting³², as it is often difficult (if not impossible) to map reports in literature back to the actual template or atlas employed in the report. Most prominently, the ubiquitous reporting of results (e.g., peak coordinates) “*in the MNI (Montreal Neurological Institute) standard space*” is inadequate because there is a large portfolio of templates^{6,33,34,23} developed and distributed by the MNI. For example, the widespread FSL toolbox³⁵ references its results to an “*MNI space*” neither created nor officially distributed by the MNI.

To address the need for a centralized resource for the archiving and redistribution of templates and atlases that allows programmatic access to human and nonhuman brain templates³⁶, we have developed TEMPLATEFlow (Figure 1). The resource is envisioned to support the emergence of processing and analysis workflows^{37,38} that brain mapping is witnessing, while addressing the above concerns threatening reproducibility. TEMPLATEFlow is modular, and both data and software are version-controlled. The resource allows researchers to use templates “off-the-shelf” and share new ones. An online documentation hub provides further details and facilitates use (<http://www.templateflow.org>).



To investigate the heterogeneous use of MNI templates in the neuroimaging literature, we performed an exploratory text mining analysis. Across the entire corpus of articles published in two leading methodological journals (*NeuroImage* and *NeuroImage: Clinical*), we identified 6,048 articles containing 14,870 sentences that included the term MNI. After preprocessing the sentence text, we used latent Dirichlet allocation⁴² to create a topic model of the surveyed articles. A qualitative inspection of topic-word loadings suggested that of the 12 topics we identified, three provided insight into the likely provenance of the templates that an article used (Figure 2). Two of these corresponded to two software packages widely used in neuroimaging—SPM (topic 6) and FSL (topic 10)—each of which is distributed with particular versions of the MNI template. The third (topic 4) related to the International Consortium for Brain Mapping (ICBM) and the MNI itself, the institutions that oversee the creation and curation of the MNI template portfolio.

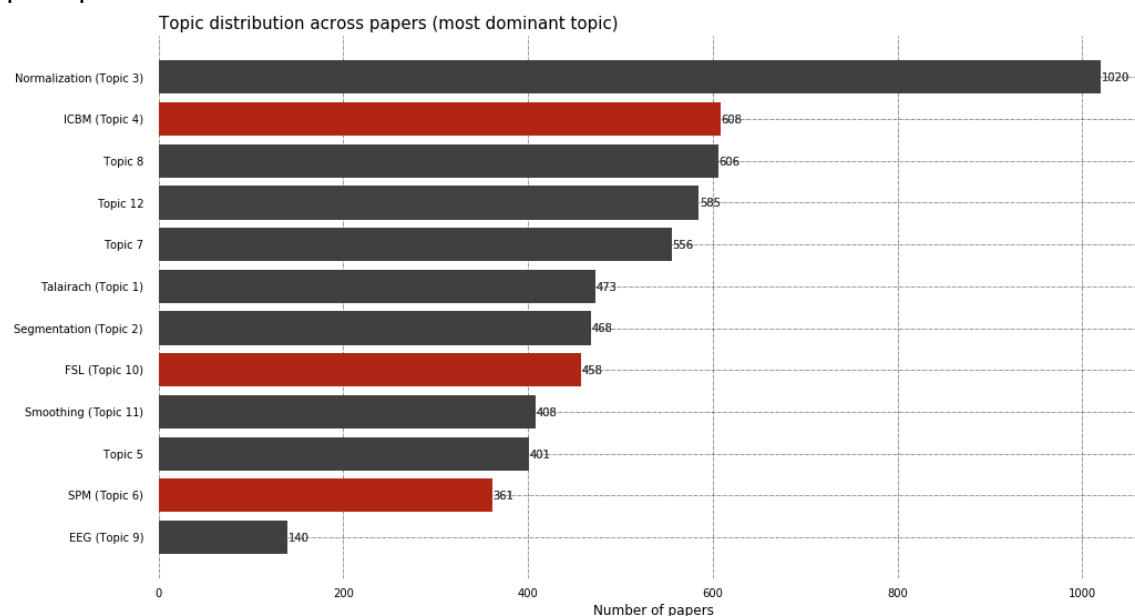


Figure 3. Topic distribution across articles. Each bar indicates the number of surveyed articles in which a topic was identified as dominant. The dominant MNI-related topic in a plurality of surveyed articles was related to general normalisation and registration. Topics corresponding to template provenance (topic 4: ICBM, topic 6: SPM, and topic 10: FSL) were each dominantly present in hundreds of articles, underscoring the heterogeneity of MNI templates used in the literature. The absence of provenance-related terms in many other articles showcases the potential difficulties in determining the template resources that were used in a neuroimaging study.

To demonstrate the heterogeneity of MNI template usage, we sorted articles according to their dominant topic (i.e., the topic with the highest model score in MNI-related sentences; Figure 3). We found that hundreds of articles featured each of the three provenance-related topics we had previously identified, underscoring that “MNI template” can refer to any of a family of templates and is not a unique identifier. As a matter of fact, studies carried out with SPM96⁴³ and earlier versions report their results “in MNI space” with reference to the single-subject Colin 27 average template⁴⁴. However, beginning with SPM99, SPM updated its definition of “MNI

space” to the template that MNI released in 2001: an average of 152 subjects from the ICBM database, aligned by means of linear registration. In SPM12 (the latest release at the time of writing), the meaning of “MNI” varies by submodule: different modules alternately use the Linear MNI152 template and a new, nonlinear revision from 2009. By contrast, the “MNI space” template bundled with the FSL neuroimaging toolbox³⁵ was developed by Dr. A. Janke in collaboration with MNI researchers. Although it was generated under the guidance of and using the techniques of the 2006 release of nonlinear MNI templates, this template is not in fact part of the official portfolio distributed by MNI. Nonetheless, our results suggest that the MNI templates bundled with SPM and FSL have historically gained broader currency as a result of the widespread use of these software libraries.

Although our results are only a first approximation, and although they do not provide insight into the provenance of the majority of MNI templates, they present several important cases for consideration. First, absent an unambiguous reporting nomenclature (such as Research Resource Identifiers [RRIDs]⁴⁵), the widespread use of ambiguous terms like “MNI” presents a potential barrier for reproducing results and increases the chance of misapplying coordinates or references from an incorrect space. Second, absent a readily accessible centralised repository, researchers might often default to templates that are easy to access, many of which are tied to specific software packages. Third, as illustrated by the changing definition of “MNI” in SPM, template references can change as new technologies emerge, suggesting an essential need for version control systems. Finally, our analysis illustrates a fourth requirement for unambiguous reporting of results: a consensus regarding the minimally sufficient provenance-related information to report in studies and to distribute with templates. We sought to develop TEMPLATEFLOW with these features in mind.

A version-controlled archive of neuroimaging templates maximizing the accuracy in reporting spatially standardized results.

TEMPLATEFLOW is a cloud-based repository of human and nonhuman imaging templates paired with a Python-based client for programmatically accessing template resources. TEMPLATEFLOW addresses the need for a standard, centralised repository of templates and corresponding atlases and metadata. The TEMPLATEFLOW ARCHIVE has a tree-directory structure, metadata files, and data files following a standard inspired by the Brain Imaging Data Structure (BIDS)⁴⁶. BIDS is a widespread standard that balances the needs for human- and machine-readability. BIDS prescribes a file naming scheme comprising a series of key-value pairs (called “entities”) that are ordered hierarchically.

The most salient entity is the template identifier (signified with the key `tpl-`), whose value is an alphanumeric label that is unique across the ARCHIVE (e.g., `tpl-MNI152Lin`). Table 1 enumerates the 15 templates currently distributed with the ARCHIVE, and their corresponding unique identifiers. The unique identifier resolves the issue of inaccurate reporting, as it unambiguously designates one specific template. In addition, because the repository is versioned, researchers can easily retrieve and report the exact version of the template or atlas that was used in their study. Supplementary Table S1 summarizes the available entities and shows a segment of the file organization of the ARCHIVE. For each template, the TEMPLATEFLOW database includes one or more reference volumetric template images (e.g., one T1-weighted and one T2-weighted average map; all must be in register), a set of atlas labels and voxelwise

annotations defined with reference to the template image, and additional files containing the template and atlas metadata. Figure 4 summarizes the data types and metadata that can be stored in the ARCHIVE. Figure 5 provides an overview of the ARCHIVE's metadata specification.

Cloud storage for the ARCHIVE is supported by the Open Science Framework (osf.io) and Amazon's Simple Storage Service (S3). Version control, replication, and synchronisation of template resources across filesystems is managed with DATA Lad⁴⁷.

Table 1. Digital templates included in TEMPLATEFLOW. *TEMPLATEFLOW is designed to maximise the discoverability and accessibility of new templates, minimise redundancies in template creation, and promote standardisation of processing workflows. To enhance visibility of existing templates, TEMPLATEFLOW includes a web-based browser indexing all files in the TEMPLATEFLOW ARCHIVE (<https://www.templateflow.org/usage/archive/>).*

Template (Unique ID)	Description
MNI152Lin ^{33,39,40}	Neurotypical adult human template created as the average from a linear mapping of 152 subjects from the MNI cohort registered to the earlier MNI305 template.
MNI152NLin2009cAsym ^{34,48}	Neurotypical adult human template created from an unbiased nonlinear averaging of 152 subjects. The mapping and averaging proceeded over 40 iterations beginning from the earlier MNI152 template.
MNI152NLin2009cSym ^{34,48}	Symmetrised version of MNI152NLin2009cAsym with identical data in left and right hemispheres.
MNI152NLin6Asym ²³	FSL's version of the MNI152 neurotypical adult human template created using iterative nonlinear registration and averaging.
MNI152NLin6Sym ²³	Symmetrised version of MNI152NLin6Asym with identical data in left and right hemispheres.
MNIInfant ³⁴	Series of human infant templates created from 11 cohorts of infants and young children. Each cohort spans a different age range between 0 and 60 months.
MNIPediatricAsym ^{34,48}	Series of human pediatric templates created from 6 partially overlapping cohorts of children and young adults. Each cohort spans a different age range between 4.5 and 18.5 years.
NKI ^{14,49}	Template created for the NKI-Rockland sample using ANTs diffeomorphic registration and averaging.
OASIS30ANTs ^{14,50}	Template created using ANTs diffeomorphic registration and averaging for the Open Access Series of Imaging Studies (OASIS).
PNC ⁴⁹	Pediatric and young adult template created using ANTs diffeomorphic registration and averaging for the

Philadelphia Neurodevelopmental Cohort.	
UNCInfant ⁵¹	Series of human infant templates created from a 95-subject longitudinal sample comprising three scans: as neonates, as one-year-olds, and as two-year-olds.
WHS ^{52–55}	Waxholm space template created as an atlas of the Sprague-Dawley rat brain.
Fischer344 ⁵⁶	Rat template created as the average of 41 four-month-old animals from the Fischer 344 strain.
fsLR ⁵⁴	Surface-based Freesurfer template created for the Human Connectome Project (HCP).
fsaverage ⁵⁷	Surface-based average Freesurfer template.

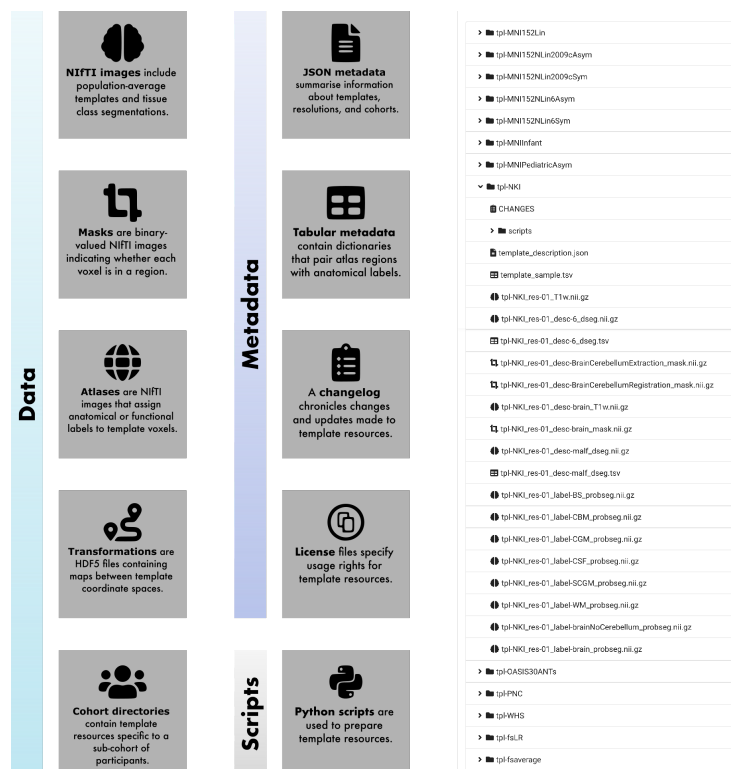


Figure 4. The *TEMPLATEFLOW ARCHIVE* contains template resources. Left, common file formats included in the *TEMPLATEFLOW ARCHIVE*. Right, view of the *TEMPLATEFLOW ARCHIVE* browser, accessible from templateflow.org, with a single template resource directory expanded. Template data are archived using a BIDS-like directory structure, with top-level directories for each template. Each directory contains image files, annotations, and metadata for that template. Following BIDS specifications, volumetric data are stored in NiftI format and surface data with the GIFTI2 format.

A key neuroimaging resource developed with the best software engineering standards and easily operable by machines.

TEMPLATEFLOW's Python client provides human users and software tools such as *fMRIPrep* with reliable and programmatic access to the archive. The client can be integrated seamlessly into image processing workflows to handle requests for template resources on the fly. It features an intuitive application programming interface (API) that can query the TEMPLATEFLOW ARCHIVE for specific files (Figure 6). The BIDS-inspired organization enables easy integration of tools and infrastructure designed for BIDS (e.g., the Python client uses PyBIDS⁵⁸ to implement the queries

listed in Table S1). To query TEMPLATEFLOW, a user can submit a list of arguments corresponding to the BIDS-like key-value pairs in each entity's file name (e.g., `atlas=Schaefer2018` to return files containing voxelwise annotations from the 2018 Schaefer atlas⁵⁹).

To integrate template resources into neuroimaging workflows, traditional approaches required deploying an oftentimes voluminous tree of prepackaged data to the filesystem. By contrast, the TEMPLATEFLOW client implements *lazy loading*, which permits the base installation to be extremely lightweight. Instead of distributing neuroimaging data with the installation, TEMPLATEFLOW allows the user to dynamically pull from the cloud-based storage only those resources they need, as they need them. After a resource has been requested once, it remains cached in the filesystem for future utilization.

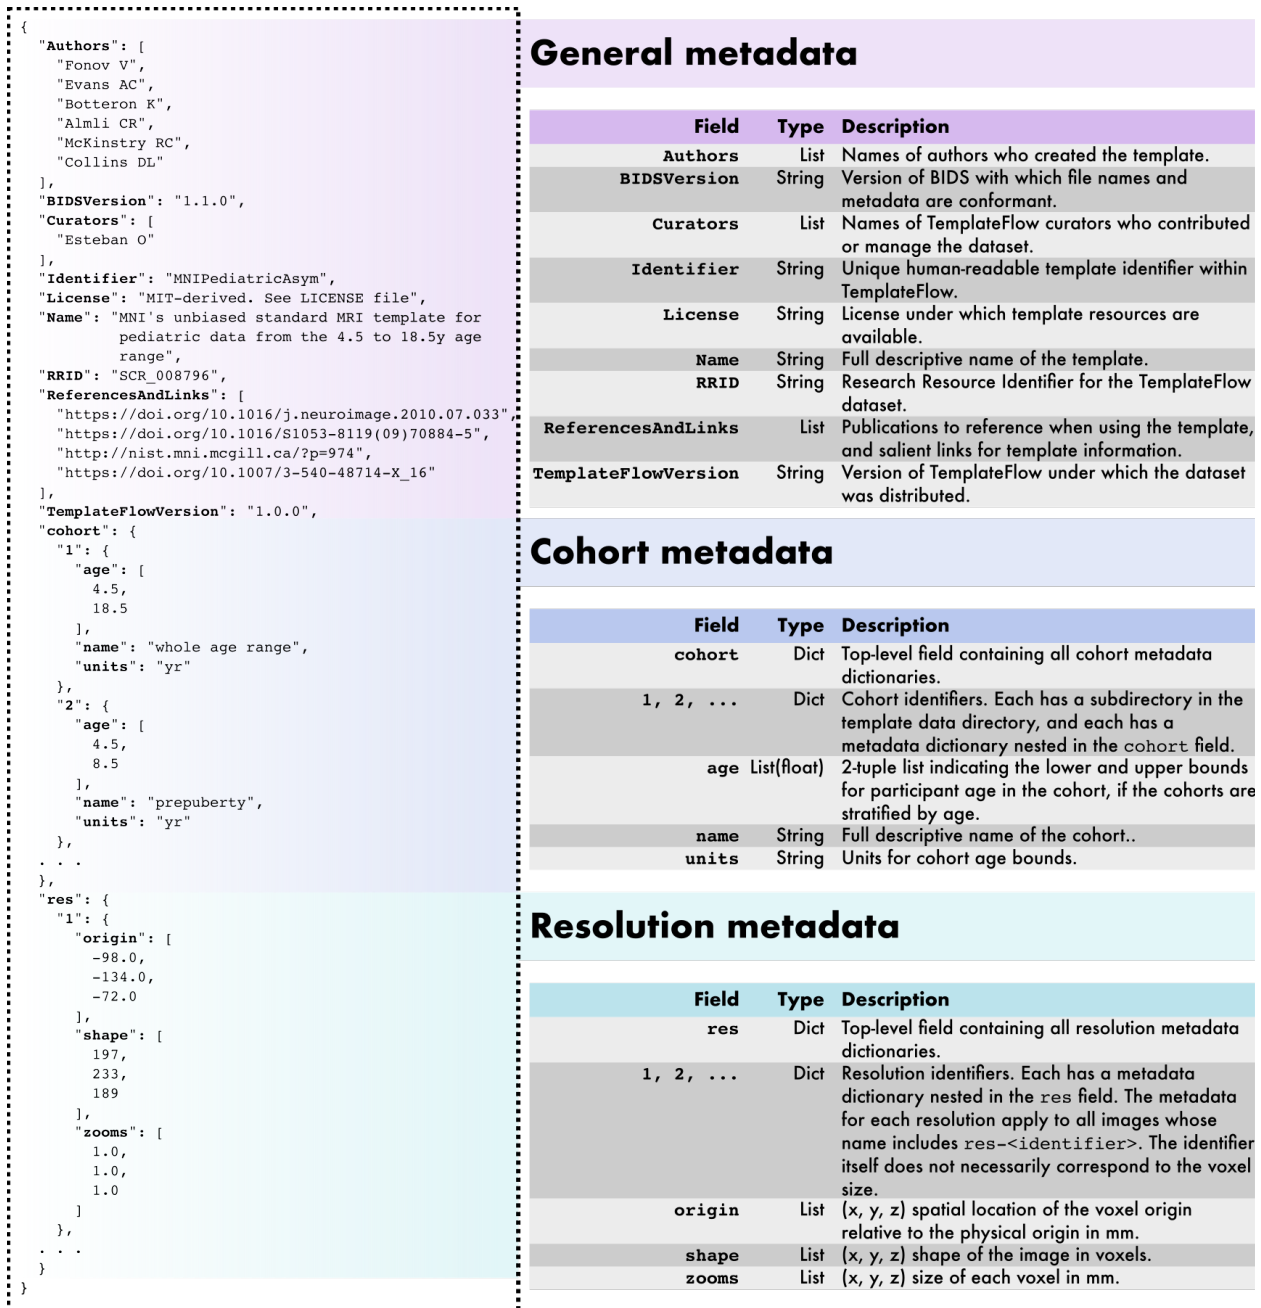


Figure 5. Overview of *TEMPLATEFLOW* metadata specification. *TEMPLATEFLOW* metadata are formatted as JavaScript Object Notation (JSON) files located in each *TEMPLATEFLOW* dataset. An example *template_description.json* metadata file is displayed at left (for the pediatric MNI template). In addition to general template metadata, *TEMPLATEFLOW* datasets can contain cohort-level and resolution-level metadata, which are nested within the main metadata dictionary and apply only to subsets of images in the dataset.

We demonstrate benefits of centralising templates in general, and the validity of the *TEMPLATEFLOW* framework in particular, via its integration into *fMRIPrep*³⁸, a functional MRI preprocessing tool. This integration provides *fMRIPrep* users with flexibility to spatially

normalize their data to any template available in the ARCHIVE. This integration has also enabled the development of *fMRIPrep* adaptations, for instance to pediatric populations or rodent imaging, utilizing suitable templates from the archive. The uniform interface provided by the BIDS-like directory organisation and metadata enables straightforward integration of new templates into workflows equipped to use TEMPLATEFLOW templates.

TEMPLATEFLOW also makes use of standards of practice from the software engineering industry, leveraging continuous delivery (CD) and continuous integration (CI) tools to automate backup and synchronisation of data across projects in the `templateflow` organisation on GitHub. CI and CD keep the web-based archive browser up to date by automatically indexing data files.

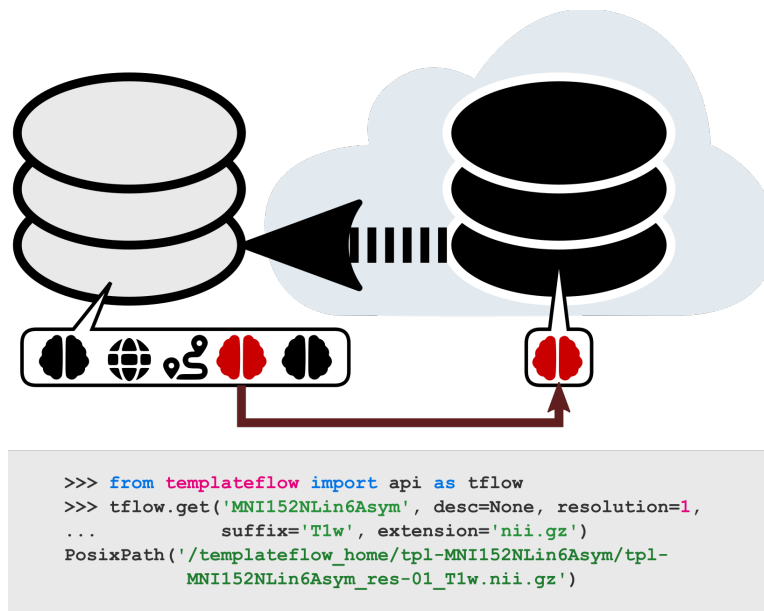


Figure 6. Example usage of the Python-based TEMPLATEFLOW client. After importing the API, the user submits a query for the T1-weighted FSL version of the MNI template at 1 mm resolution. The client first filters through the archive, identifies any files that match the query, and finds their counterparts in cloud storage. It then downloads the requested files and returns their paths in the local TEMPLATEFLOW installation directory. Future queries for the same resource can be completed without any re-downloading.

Community-driven, peer-reviewed contribution process.

A centralised repository for neuroimaging templates should also address the needs of template creators, enabling peer-reviewed integration of new templates with minimal informatic overhead. Inspired by the Conda-forge community repository and the Journal of Open Source Software (JOSS), the GitHub-based TEMPLATEFLOW organisation is a site for dialogue between members of the neuroimaging community and TEMPLATEFLOW ARCHIVE curators. GitHub issues offer any community member the ability to share their needs with developers and ARCHIVE curators, for instance by identifying templates or workflow features for potential inclusion in the

project. Pull requests provide a means for members of the community to directly contribute code or template resources to the TEMPLATEFLOW ARCHIVE.

This peer-reviewed contribution process is facilitated through the Python-based TEMPLATEFLOW MANAGER. The TEMPLATEFLOW MANAGER automates the work of synchronising data from a local directory to cloud storage in OSF. Furthermore, it creates a GitHub repository containing `git-annex` pointers that enable DATALAD to download template data from cloud storage to any machine with a copy of the repository. Finally, it opens a new pull request to propose adding the newly contributed template repository into the main TEMPLATEFLOW ARCHIVE (Figure 7).

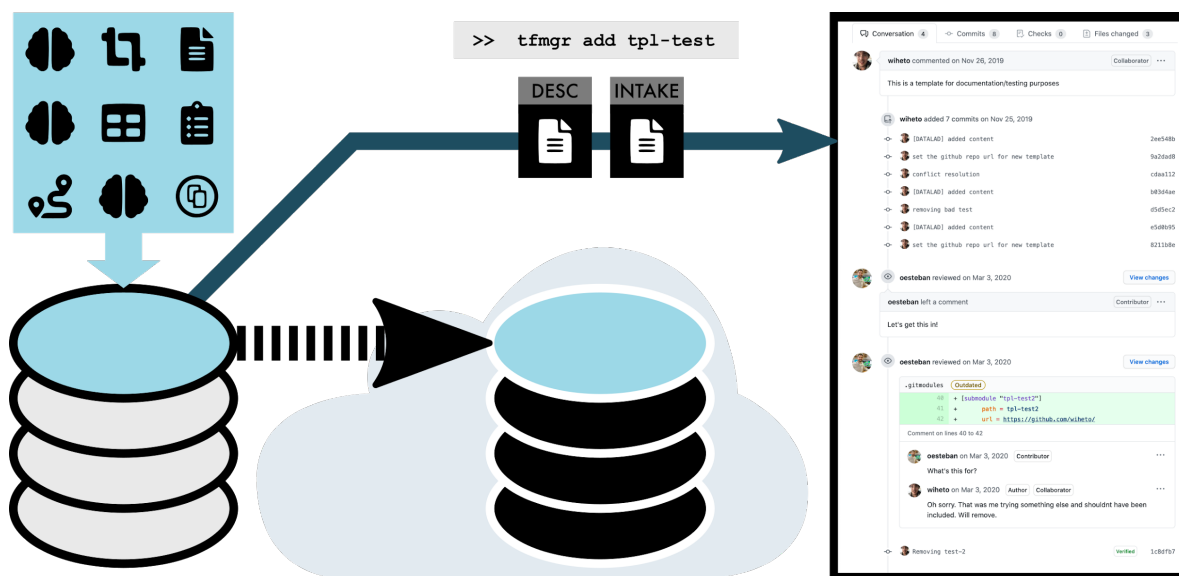


Figure 7. To contribute a new template to TEMPLATEFLOW, members of the community first organise template resources to conform to the BIDS-like TEMPLATEFLOW structure. Next, `tfmgr` synchronises the resources to OSF cloud storage and opens a new pull request proposing the addition of the new template. A subsequent peer review process ensures that all data are conformant with the TEMPLATEFLOW standard. Finally, TEMPLATEFLOW curators merge the pull request, thereby adding the template into the archive.

Discussion

The use of templates in neuroimaging is ubiquitous, and the emerging challenges regarding template use accordingly merit immediate attention. In an early perspective, Van Essen identified a set of desiderata for brain templates³⁶. Above and beyond anatomical fidelity, he called for connecting templates in a “*federation of databases*” with “*powerful and flexible options for searching, selecting, and visualizing data*”. Finally, he stressed the importance of resource accessibility. TEMPLATEFLOW provides a clear foundation for a framework that satisfies all of the aforementioned desiderata. Furthermore, the ARCHIVE supports unambiguous identification of resources, a programmatically extensible platform for interfacing with template data, and a starting point for future investigations of inter-template concordance and robustness.

TEMPLATEFlow addresses practical issues biasing current template usage and creation.

When researchers develop a new brain template or atlas for public dissemination, there exists no standard channel or format for distributing their work. With no central repository or uniform organisational scheme, template creators are often tasked with the responsibilities of maintaining template resources and managing access, sometimes on an *ad hoc* basis. The work resulting in a template is reviewed only at publication time; subsequent template updates can go unreviewed, and any academic consensus that emerges after publication might not be associated with the original template resources at all.

Conversely, users are confronted with a surfeit of available templates and atlases, many with unclear provenance, and with the attendant challenges of locating them and integrating them into workflows. As illustrated by our text mining analysis, software libraries can further bias template usage, and the lack of consistent nomenclature between packages (as is the case for MNI templates) introduces further ambiguity in scientific communication. Leading MRI analysis libraries, including FSL and SPM, are packaged with a limited set of default template resources derived from a neurotypical adult human population. For the many researchers who study phenomena in other populations and species, these resources can be inadequate, leading researchers to *ad hoc* decisions regarding the creation and use of templates.

Together, user-facing and developer-facing issues contribute to an environment of ambiguity. Even where best practices are known, they can be difficult to locate and follow. Awareness and discoverability of extant template resources is limited by access hurdles and the use of private, decentralised communication channels, thereby driving unnecessary proliferation of template options. Integrating a template into software requires a custom solution for every new template, increasing the burden on developers. By aggregating resources in a centralised and freely accessible archive, TEMPLATEFlow maximises exposure of the scientific community to new templates and facilitates the dissemination of new template-based knowledge.

A comprehensive, eclectic, and unambiguous template portfolio. Using a software package's default template without awareness of alternatives can yield suboptimal results with adverse effects on both reproducibility and reporting. To mitigate these difficulties, TEMPLATEFlow aggregates representative template data for populations and species of scientific interest in a centralised, community-run archive. Furthermore, TEMPLATEFlow's organisational scheme enables univocal identification of brain templates, resolving the ambiguity of references to "MNI" space and thereby facilitating experimental reproducibility. More broadly, TEMPLATEFlow defines a BIDS-like language for references to brain templates and atlases⁴⁶, endowing researchers with a common vocabulary so that the results they report in future work can be traced unambiguously to specific resources and replicated precisely. The archive provides template data across a range of granularities—for instance, researchers who study developmental trajectories can use a generally representative youth template or a template specific to a narrower age range according to the objectives of their research.

An improved platform for vetting template-based knowledge and building consensus.

GitHub's pull request system and integrated peer review process provide a public forum for discussion and vetting of newly proposed templates and annotations. Each template can be

treated as a versioned living document that is continually reviewed and republished as necessary. If a user questions the validity or currentness of archive resources, the platform provides an immediate channel for publicly raising and addressing concerns. The discussion and vetting process forms a record of outstanding issues and concerns with existing resources that can inform researchers about the strengths and limitations of resources available to them. Ultimately, the centralisation of resources and discussion can help create a scientific consensus. As new technology enables the refinement of template resources, TEMPLATEFLOW's integration with GitHub's version control system provides a built-in way to update existing templates and track revisions.

An intuitive API minimises the window for human error. The TEMPLATEFLOW client implements an intuitive query system that enables programmatic retrieval of template resources from the archive. Using a command-line tool or the Python Client API, scientists can easily integrate TEMPLATEFLOW into their neuroimaging analysis workflows to automate access to and use of template resources. Automation of workflows further promotes reproducibility by removing potential points of inconsistent or erroneous usage. Additionally, the BIDS-like organisation scheme of TEMPLATEFLOW resources facilitates integration with BIDS apps. Within *fMRIPrep*, for instance, the TEMPLATEFLOW API enables flexible and systematic normalisation of preprocessed images to any template space requested by the user.

Delineating “standardness” of templates and coordinate spaces. A dispute over what qualifies as a “standard” space stems in part from the limitations of a “one-size-fits-all” approach to coordinate spaces, in part from differing study objectives, and in part from the limitations of knowledge about the organisation of the brain. In particular, the substantial variability between brains manifests not only in the morphological positions of brain features, but also in the failure of some landmarks to consistently manifest in all subjects. Functionally, deep phenotyping studies⁶⁰ have similarly revealed that the subnetwork structure of individual human brains exhibits features not present in population-level averages, up to and including entire large-scale functional networks⁶¹. Furthermore, the relationships between many brain structures and their functions either are imprecisely characterised⁶², defy intuitive ontologies^{63,64}, or vary among individual subjects. Such considerations call into question the “standardness” of standard spaces derived from population averages.

These challenges are doubly amplified when researchers aim to identify inter-species homologies in the architecture of the nervous system, an endeavour that might ultimately require defining an abstract relational space that leverages spatial geometry^{65,66} rather than one tethered to an explicit, population-average template. The scope of the present work does not immediately encompass such challenges, but it is hoped that the resource introduced here can provide a starting point for future work in these veins. As registration frameworks and standard space definitions expand beyond anatomical averages and incorporate information from additional modalities^{67,68}, we intend TEMPLATEFLOW to grow to accommodate these new types of resources.

Online Methods

MNI space text mining analysis. To investigate the use of the term “MNI” in the neuroimaging literature, we conducted an exploratory text mining analysis. For this, we used the Elsevier API to download the entire corpus of two leading journals of neuroimaging methodology, *NeuroImage* and *NeuroImage: Clinical*. In this way, we retrieved a total 16,812 full-text articles that were subsequently tokenized into lists of sentences. A scan of these sentences revealed 14,870 sentences across 6,048 articles that contained the word “MNI”. Sentences were cleaned (i.e., removing punctuation, single letters, accents, numbers) and tokenized into words, which were subsequently lemmatized (i.e., converted to base form) using the NLTK wordnet lemmatizer. From the lemmatized words, we filtered out stopwords (i.e., NLTK stopwords and a custom list) and included words with a frequency above 10 as part of our “dictionary”; this yielded a dictionary size of 2,340 words.

Next, we computed a sparse dictionary by article count matrix (i.e., 2,340 x 6,048), on which we performed topic modelling with latent Dirichlet allocation⁴² (LDA; implementation from scikit-learn with the learning decay hyperparameter set to 0.7). The number of topics (k=12) was selected by identifying the LDA model yielding the lowest perplexity (Blei, Ng & Jordan 2003). The 20 words from the dictionary that loaded the highest on the 12 topics were visualized using word clouds. In addition, for each article we identified the most dominant topic and plotted the distribution of topics across articles.

Design and architecture. TEMPLATEFLOW comprises 4 cardinal components: (i) a cloud-based archive, (ii) a Python client for programmatically querying the archive, (iii) automated systems for synchronising and updating archive data, and (iv) inter-template registration workflows. Here, we discuss the details of each component’s implementation in turn, as well as the manner in which they interact with one another to form a cohesive whole.

The TEMPLATEFLOW ARCHIVE. The archive itself comprises directories of template data in cloud storage. The data are stored on Google Cloud using the Open Science Framework (OSF) and on Amazon’s Simple Storage Service (S3). Prior to storage, all template data must be named and organised in directories conformant to a data structure inspired by and compatible with the Brain Imaging Data Structure (BIDS) standard⁴⁶. The precise implementation of this data structure is a living document and is detailed on the TEMPLATEFLOW homepage (<http://www.templateflow.org>). We detail several critical features here.

The archive is organised hierarchically, and descriptive metadata follow a principle of inheritance: any metadata that apply to a particular level of the archive also apply to all inferior levels. At the top level of the hierarchy are directories corresponding to each archived template. If applicable, within each template directory are directories corresponding to sub-cohort templates. Names of directories and resource files constitute a hierarchically ordered series of key-value pairs terminated by a suffix denoting the datatype. For instance, `tpl-MNIPediatricAsym_cohort-3_res-high_T1w.nii.gz` denotes a T1-weighted template image file for resolution “high” of cohort 3 in the `MNIPediatricAsym` template (where the definitions of each resolution and cohort are specified in the template metadata file, `template_description.json`). The most common TEMPLATEFLOW datatypes are indexed in Table 1 of the main manuscript; an exhaustive list is available in the most current version of the BIDS standard (<https://bids.neuroimaging.io/>).

Within each directory, template resources include image data, atlas and template metadata, transform files, licenses, and curation scripts. All image data are stored in gzipped NIfTI format and are conformed to RAS+ orientation (i.e., left-to-right, posterior-to-anterior, inferior-to-superior, with the affine qform and sform matrices corresponding to a cardinal basis scaled to the resolution of the image). Template metadata are stored in a JavaScript Object Notation (JSON) file called `template_description.json`; an overview of metadata specifications is provided in Figure 5 of the main manuscript. In brief, template metadata files contain general template metadata (e.g., authors and curators, references), cohort-specific metadata (e.g., ages of subjects included in each cohort), and resolution-specific metadata (e.g., dimensions of images associated with each resolution). Atlas metadata are often stored in TSV format and specify the region name corresponding to each atlas label. Transform files are stored in HDF5 format and are generated as a diffeomorphic composition of ITK-formatted transforms mapping between each pair of templates.

The archive has a number of client-facing access points to facilitate browsing of resources. Key among these is the archive browser on the TEMPLATEFLOW homepage, which indexes all archived resources and provides a means for researchers to take inventory of possible templates to use for their study.

The Python client. TEMPLATEFLOW is distributed with a Python client that can submit queries to the archive and download any resources as they are requested by a user or program. Valid query options correspond approximately to BIDS key-value pairs and datatypes. A compendium of common query arguments is provided in Table 1 of the main manuscript, and comprehensive documentation is available on the TEMPLATEFLOW homepage.

When a query is submitted to the TEMPLATEFLOW client, the client begins by identifying any files in the archive that match the query. To do so, it uses PyBIDS⁵⁸, which exploits the BIDS-like architecture of the TEMPLATEFLOW ARCHIVE to efficiently scan all directories and filter any matching files. Next, the client assesses whether queried files exist as data in local storage. When a user locally installs TEMPLATEFLOW, the local installation initially contains only lightweight pointers to files in OSF cloud storage. These pointers are implemented using DATAHAD, a data management tool that extends git and git-annex⁴⁷. TEMPLATEFLOW uses DATAHAD principally to synchronise datasets across machines and to perform version control by tracking updates made to a dataset.

If the queried files are not yet synchronised locally (i.e., they exist only as pointers to their counterparts in the cloud), the client instructs DATAHAD to retrieve them from cloud storage. In the event that DATAHAD fails or returns an error, the client falls back on redundancy in storage and downloads the file directly from Amazon's S3. When the client is next queried for the same file, it will detect that the file has already been cached in the local filesystem. The use of resource pointers with the client thus enables lazy loading of template resources. Finally, the client confirms that the file has been downloaded successfully. If the client detects a successful download, it returns the result of the query; in the event that it detects a synchronisation failure, it displays a warning for each queried file that encountered a failure.

Continued functionality and operability of the client is ensured through an emphasis on maximising code coverage with unit tests. Updating the client requires successful completion of all unit tests, which are automatically executed by continuous integration (CI) and continuous delivery (CD) services connected to GitHub.

Ancillary and managerial systems. TEMPLATEFLOW includes a number of additional systems and programs that serve to automate stages of the archive update process, for instance addition of a new template or revision of current template resources. To facilitate the update and extension process, TEMPLATEFLOW uses GitHub actions to automatically synchronise dataset information so that all references remain up-to-date with the current dataset. These actions are triggered whenever a pull request to TEMPLATEFLOW is accepted. For example, GitHub actions are used to update the TEMPLATEFLOW archive browser so that it displays all template resources as they are uploaded to the archive.

Whereas the TEMPLATEFLOW client synchronises data from cloud storage to the local filesystem, a complementary TEMPLATEFLOW manager handles the automated synchronisation of data from the local filesystem to cloud storage. The Python-based manager is also used for template intake, i.e., to propose the addition of new templates to the archive. To propose adding a new template, a user first runs the TEMPLATEFLOW manager using the `tfmgr add <template_id> --osf-project <project_id>` command.

The manager begins by using the TEMPLATEFLOW client to query the archive and verify that the proposed template does not already exist. After verifying that the proposed template is new, the manager synchronises all specified template resources to OSF cloud storage. It then creates a fork of the `tpl-intake` branch of the TEMPLATEFLOW GitHub repository and generates an *intake* file in Tom's Obvious Minimal Language (TOML) markup format; this intake file contains a reference to the OSF project where the manager has stored template resources. The TEMPLATEFLOW MANAGER commits the TOML intake file to the fork and pushes to the user's GitHub account. Finally, it retrieves template metadata from `template_description.json` and uses the metadata to compose a pull request on the `tpl-intake` branch. This pull request provides a venue for discussion and vetting of the proposed addition of a new template.

Inter-template registration workflow. To effect the flow of knowledge across template spaces, TEMPLATEFLOW includes a workflow for computing robust transformations between any pair of adult human template spaces. To compute a transformation between two template spaces, the inter-template registration workflow makes use of 10 of the high-quality T1-weighted adult human brain images used in the creation of the MNI 152 template portfolio. In the first step of the workflow, these 10 images are registered to both template spaces using the symmetric normalization (SyN) algorithm²⁴. Next, a 10-channel registration is performed in ANTs using the SyN algorithm. Thus, the workflow computes a single transformation that simultaneously optimises the alignment between all 10 images in both coordinate spaces.

References

1. Lüders, E., Steinmetz, H. & Jäncke, L. Brain size and grey matter volume in the healthy human brain: *NeuroReport* **13**, 2371–2374 (2002).
2. Tosun, D., Siddarth, P., Levitt, J. & Caplan, R. Cortical thickness and sulcal depth: insights on development and psychopathology in paediatric epilepsy. *BJPsych Open* **1**, 129–135 (2015).
3. Tavor, I. *et al.* Task-free MRI predicts individual differences in brain activity during task performance. *Science* **352**, 216–220 (2016).

4. Mars, R. B., Passingham, R. E. & Jbabdi, S. Connectivity Fingerprints: From Areal Descriptions to Abstract Spaces. *Trends Cogn. Sci.* **22**, 1026–1037 (2018).
5. Sowell, E. R. *et al.* Mapping cortical change across the human life span. *Nat. Neurosci.* **6**, 309–315 (2003).
6. Good, C. D. *et al.* A Voxel-Based Morphometric Study of Ageing in 465 Normal Adult Human Brains. *NeuroImage* **14**, 21–36 (2001).
7. Courchesne, E. *et al.* Normal Brain Development and Aging: Quantitative Analysis at in Vivo MR Imaging in Healthy Volunteers. *Radiology* **216**, 672–682 (2000).
8. Brett, M., Johnsrude, I. S. & Owen, A. M. The problem of functional localization in the human brain. *Nat. Rev. Neurosci.* **3**, 243–249 (2002).
9. von Economo, C. F., Koskinas, G. N. & Triarhou, L. C. *Atlas of cytoarchitectonics of the adult human cerebral cortex*. vol. 10 (Karger Basel, 2008).
10. Brodmann, K. *Beiträge zur histologischen Lokalisation der Grosshirnrinde: Fünfte Mitteilung: Über den allgemeinen Bauplan des Cortex Pallii bei den Mammaliern und zwei homologe Rindenfelder im besonderen. Zugleich ein Beitrag zur Furchenlehre*. vol. 6 (Suppl.) (J. Psychol. Neurol., 1906).
11. Schurr, P. H. & Merrington, W. R. The Horsley–Clarke stereotaxic apparatus. *Br. J. Surg.* **65**, 33–36 (1978).
12. Talairach, J., David, M., Tournoux, P., Corredor, H. & Kvasina, J. *Atlas d'anatomie stéréotaxique: repérage radiologique indirect des noyaux gris centraux des régions mésencéphalo-sous-optique et hypothalamique de l'homme*. (Masson, 1957).
13. Talairach, J. & Tournoux, P. *Co-planar stereotaxic atlas of the human brain*. (Georg Thieme Verlag/Thieme Medical Publishers, 1988).
14. Avants, B. B. *et al.* A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage* **54**, 2033–2044 (2011).
15. Paxinos, G. & Watson, C. *The rat brain in stereotaxic coordinates: hard cover edition*. (Elsevier Academic Press, 1997).
16. Martin, R. F. & Bowden, D. M. *Primate Brain Maps: Structure of the Macaque Brain*. (Elsevier, 2000).
17. Dickie, D. A., Job, D. E., Gonzalez, D. R., Shenkin, S. D. & Wardlaw, J. M. Use of Brain MRI Atlases to Determine Boundaries of Age-Related Pathology: The Importance of Statistical Method. *PLOS ONE* **10**, e0127939 (2015).
18. Matsuzawa, J. *et al.* Age-related Volumetric Changes of Brain Gray and White Matter in Healthy Infants and Children. *Cereb. Cortex* **11**, 335–342 (2001).
19. Fonov, V. *et al.* Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* **54**, 313–327 (2011).
20. Buckner, R. L. *et al.* A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization:

- reliability and validation against manual measurement of total intracranial volume. *NeuroImage* **23**, 724–738 (2004).
21. Calabrese, E., Badea, A., Watson, C. & Johnson, G. A. A quantitative magnetic resonance histology atlas of postnatal rat brain development with regional estimates of growth and variability. *NeuroImage* **71**, 196–206 (2013).
22. Szulc, K. U. *et al.* 4D MEMRI atlas of neonatal FVB/N mouse brain development. *NeuroImage* **118**, 49–62 (2015).
23. Evans, A. C., Janke, A. L., Collins, D. L. & Baillet, S. Brain templates and atlases. *NeuroImage* **62**, 911–922 (2012).
24. Avants, B., Epstein, C., Grossman, M. & Gee, J. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**, 26–41 (2008).
25. Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J. & Dalca, A. V. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Trans. Med. Imaging* **38**, 1788–1800 (2019).
26. Yoon, U., Fonov, V. S., Perusse, D. & Evans, A. C. The effect of template choice on morphometric analysis of pediatric brain data. *NeuroImage* **45**, 769–777 (2009).
27. Dickie, D. A. *et al.* Whole Brain Magnetic Resonance Image Atlases: A Systematic Review of Existing Atlases and Caveats for Use in Population Imaging. *Front. Neuroinformatics* **11**, (2017).
28. Thomas Yeo, B. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
29. Botvinik-Nezer, R. *et al.* Variability in the analysis of a single neuroimaging dataset by many teams. *Nature* **582**, 84–88 (2020).
30. Carp, J. On the Plurality of (Methodological) Worlds: Estimating the Analytic Flexibility of fMRI Experiments. *Front. Neurosci.* **6**, (2012).
31. Bohland, J. W., Bokil, H., Allen, C. B. & Mitra, P. P. The Brain Atlas Concordance Problem: Quantitative Comparison of Anatomical Parcellations. *PLoS ONE* **4**, e7200 (2009).
32. Carp, J. The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage* **63**, 289–300 (2012).
33. Mazziotta, J. *et al.* A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **356**, 1293–1322 (2001).
34. Fonov, V. *et al.* Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* **54**, 313–327 (2011).
35. Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W. & Smith, S. M. FSL. *NeuroImage* **62**, 782–790 (2012).
36. Van Essen, D. C. Windows on the brain: the emerging role of atlases and databases in

- neuroscience. *Curr. Opin. Neurobiol.* **12**, 574–579 (2002).
37. Esteban, O. *et al.* MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLOS ONE* **12**, e0184661 (2017).
38. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).
39. Mazziotta, J. C., Toga, A. W., Evans, A., Fox, P. & Lancaster, J. A Probabilistic Atlas of the Human Brain: Theory and Rationale for Its Development. *NeuroImage* **2**, 89–101 (1995).
40. Mazziotta, J. *et al.* A Four-Dimensional Probabilistic Atlas of the Human Brain. *J. Am. Med. Inform. Assoc.* **8**, 401–430 (2001).
41. Evans, A. C. *et al.* 3D statistical neuroanatomical models from 305 MRI volumes. in *1993 IEEE Conference Record Nuclear Science Symposium and Medical Imaging Conference* 1813–1817 vol.3 (1993). doi:10.1109/NSSMIC.1993.373602.
42. Blei, D. M., Ng, A. Y. & Jordan, M. I. Latent dirichlet allocation. *J Mach Learn Res* **3**, 993–1022 (2003).
43. Friston, K. J. *et al.* Statistical parametric maps in functional imaging: A general linear approach. *Hum. Brain Mapp.* **2**, 189–210 (1994).
44. Holmes, C. J. *et al.* Enhancement of MR Images Using Registration for Signal Averaging: *J. Comput. Assist. Tomogr.* **22**, 324–333 (1998).
45. Bandrowski, A. E. & Martone, M. E. RRIIDs: A Simple Step toward Improving Reproducibility through Rigor and Transparency of Experimental Methods. *Neuron* **90**, 434–436 (2016).
46. Gorgolewski, K. J. *et al.* The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Sci. Data* **3**, 160044 (2016).
47. Hanke, Michael *et al.* *datalad/datalad: First release candidate for 0.13.0 (May 05, 2020)*. (Zenodo, 2020). doi:10.5281/ZENODO.3787920.
48. Collins, D. L., Zijdenbos, A. P., Baaré, W. F. C. & Evans, A. C. ANIMAL+INSECT: Improved Cortical Structure Segmentation. in *Information Processing in Medical Imaging* (eds. Kuba, A., Šámal, M. & Todd-Pokropek, A.) vol. 1613 210–223 (Springer Berlin Heidelberg, 1999).
49. Nooner, K. B. *et al.* The NKI-Rockland Sample: A Model for Accelerating the Pace of Discovery Science in Psychiatry. *Front. Neurosci.* **6**, (2012).
50. Marcus, D. S. *et al.* Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI Data in Young, Middle Aged, Nondemented, and Demented Older Adults. *J. Cogn. Neurosci.* **19**, 1498–1507 (2007).
51. Shi, F. *et al.* Infant brain atlases from neonates to 1- and 2-year-olds. *PloS One* **6**, e18746 (2011).
52. Kjonigsen, L. J., Lillehaug, S., Bjaalie, J. G., Witter, M. P. & Leergaard, T. B. Waxholm Space atlas of the rat brain hippocampal region: Three-dimensional delineations based on

- magnetic resonance and diffusion tensor imaging. *NeuroImage* **108**, 441–449 (2015).
53. Osen, K. K., Imad, J., Wennberg, A. E., Papp, E. A. & Leergaard, T. B. Waxholm Space atlas of the rat brain auditory system: Three-dimensional delineations based on structural and diffusion tensor magnetic resonance imaging. *NeuroImage* **199**, 38–56 (2019).
 54. Papp, E. A., Leergaard, T. B., Calabrese, E., Johnson, G. A. & Bjaalie, J. G. Waxholm Space atlas of the Sprague Dawley rat brain. *NeuroImage* **97**, 374–386 (2014).
 55. Papp, E. A., Leergaard, T. B., Calabrese, E., Allan Johnson, G. & Bjaalie, J. G. Addendum to “Waxholm Space atlas of the Sprague Dawley rat brain” [*NeuroImage* 97 (2014) 374–386]. *NeuroImage* **105**, 561–562 (2015).
 56. Goerzen, D. *et al.* An MRI-Derived Neuroanatomical Atlas of the Fischer 344 Rat Brain. *Sci. Rep.* **10**, 6952 (2020).
 57. Fischl, B., Sereno, M. I. & Dale, A. M. Cortical Surface-Based Analysis II: Inflation, Flattening, and a Surface-Based Coordinate System. *NeuroImage* **9**, 195–207 (1999).
 58. Yarkoni, T. *et al.* PyBIDS: Python tools for BIDS datasets. *J. Open Source Softw.* **4**, 1294 (2019).
 59. Schaefer, A. *et al.* Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2018).
 60. Poldrack, R. A. *et al.* Long-term neural and physiological phenotyping of a single human. *Nat. Commun.* **6**, 8885 (2015).
 61. Laumann, T. O. *et al.* Functional System and Areal Organization of a Highly Sampled Individual Human Brain. *Neuron* **87**, 657–670 (2015).
 62. Cabeza, R. & Nyberg, L. Imaging Cognition II: An Empirical Review of 275 PET and fMRI Studies. *J. Cogn. Neurosci.* **12**, 1–47 (2000).
 63. Poldrack, R. A. Mapping Mental Function to Brain Structure: How Can Cognitive Neuroimaging Succeed? *Perspect. Psychol. Sci.* **5**, 753–761 (2010).
 64. Barrett, L. F. & Satpute, A. B. Large-scale brain networks in affective and social neuroscience: towards an integrative functional architecture of the brain. *Curr. Opin. Neurobiol.* **23**, 361–372 (2013).
 65. Stephan, K. E. & Kötter, R. One cortex – many maps: An introduction to coordinate-independent mapping by Objective Relational Transformation (ORT). *Neurocomputing* **26–27**, 1049–1054 (1999).
 66. Huntenburg, J. M., Bazin, P.-L. & Margulies, D. S. Large-Scale Gradients in Human Cortical Organization. *Trends Cogn. Sci.* **22**, 21–31 (2018).
 67. Robinson, E. C. *et al.* MSM: A new flexible framework for Multimodal Surface Matching. *NeuroImage* **100**, 414–426 (2014).
 68. Toga, A. W., Thompson, P. M., Mori, S., Amunts, K. & Zilles, K. Towards multimodal atlases of the human brain. *Nat. Rev. Neurosci.* **7**, 952–966 (2006).

69. Esteban, O. *et al.* MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLOS ONE* **12**, e0184661 (2017).
70. Pruim, R. H. R. *et al.* ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage* **112**, 267–277 (2015).
71. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).

Supplemental Materials

Supplementary Table 1. TEMPLATEFLOW data entities.

Data entity	API query example	Description
Template	"MNI152Lin"	The template dataset to which an image or other data file belongs.
Resolution	res=1	The image resolution. Each resolution is assigned a key, which is defined in the <code>res</code> field of <code>template_description.json</code> .
Mask	desc="brain", suffix="mask"	Indicates that the image is a binary-valued annotation, where voxels labelled 1 are part of the mask.
Discrete segmentation	desc="malf", suffix="dseg"	Indicates that the image is an integer-valued annotation. Each segmentation image file (<code>.nii.gz</code> format) is paired with a dictionary of segment names (<code>.tsv</code> format).
Probabilistic segmentation	label="CSF", suffix="probseg"	Indicates that the image is a probabilistic annotation, wherein the value of each voxel indicates the probability of that voxel belonging to the specified label.
Atlas	atlas="Schaefer", desc="7Network"	The atlas to which a segmentation file belongs.
Transformation	from="MNI152Lin", suffix="xfm"	File containing a mapping between 2 stereotaxic coordinate spaces. The source space is defined in the <code>from</code> field, while the target space is defined in the <code>tpl</code> field.
Image modality	suffix="T1w"	For non-annotation brain images, the suffix indicates whether the image is T1-weighted (T1w), T2-weighted (T2w), proton density-weighted (PD), or T2*-weighted (T2star).
Template cohort	cohort=1	Subsample of a dataset used to generate an average template.

Supplementary Table 2. Command-line interface for TEMPLATEFLOW MANAGER.

Argument	Environment variable	Specifications
<code>template_id</code>		Identifier of the template. This is the value of the <code>tpl</code> field in all file names.

<code>--osf_project OSF_PROJECT</code>	The OSF project where the template data are to be stored. The project must be writable by the user account whose credentials are specified in the <code>--osf-user</code> and <code>--osf-password</code> arguments.
<code>--osf_user OSF_USERNAME</code>	Account username or identifier for OSF cloud storage.
<code>--osf_password OSF_PASSWORD</code>	Account password for OSF cloud storage.
<code>--osf_overwrite</code>	Flag that indicates that the OSF client should force the overwrite of any existing files in the OSF project that have names conflicting with those of new files.
<code>--gh-user GITHUB_USER</code>	Account username for GitHub. The user account whose credentials are provided must have a fork of the TEMPLATEFLOW repo.
<code>--gh-password GITHUB_PASSWORD</code>	Account password for GitHub.
<code>--path</code>	Path to a local directory where template resources are located. The path must either be a directory whose name is <code>tpl-<template_id></code> or contain such a directory.
<code>--nprocs</code>	Maximum number of parallel processes to run when uploading to or fetching from OSF.

Supplementary Box 1. The TEMPLATEFLOW MANAGER utility.

As a reference, we provide a brief protocol for contribution of templates stored on OSF. The requirements for using this protocol are a GitHub account, the `hub` command-line interface for GitHub (<https://github.com/github/hub>), an OSF account, the OSF Python client (downloadable from <https://github.com/osfclient/osfclient>), the TEMPLATEFLOW Python client (downloadable from <https://github.com/templateflow/python-client>), and the TEMPLATEFLOW archive manager (downloadable from <https://github.com/templateflow/python-manager>).

1. First, locally aggregate and organise template resources in a top-level `tpl` directory according to the TEMPLATEFLOW specification (Figure 2, Table 1). The `tpl` directory must use a unique name not assigned to any existing templates in the archive. It should contain template metadata in a JSON file called `template_description.json` (Figure 5).

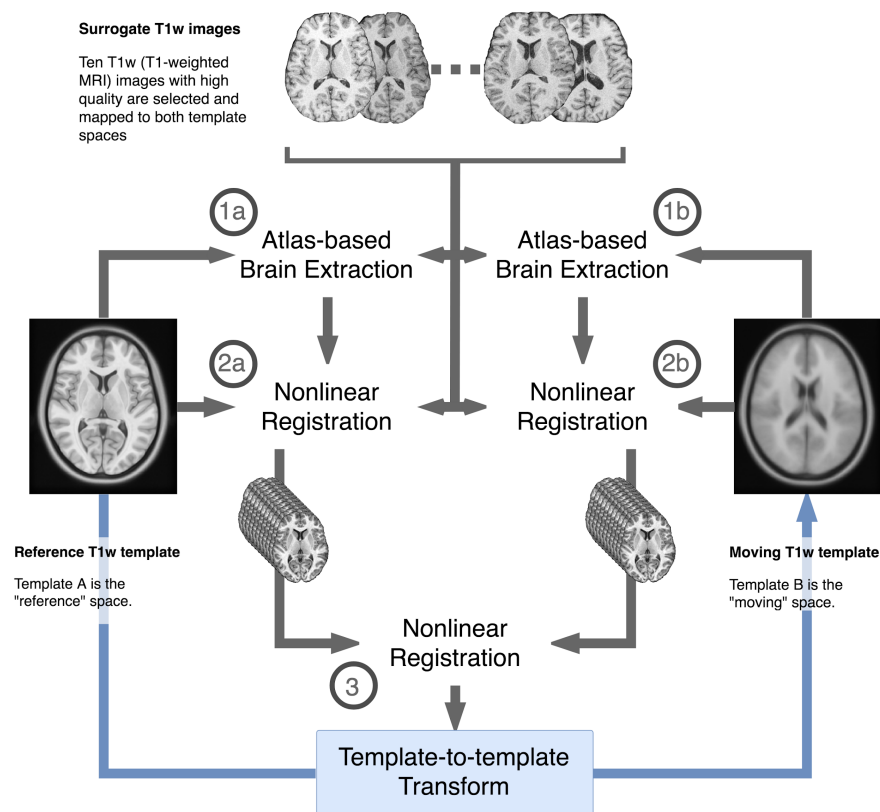
2. Next, create a new project on OSF to store templates that meet the TEMPLATEFLOW specification.
3. Execute `tfmgr`. The minimal syntax for adding a new template is `tfmgr add <template_id> --osf-project <project_id>`. Arguments to `tfmgr` are detailed in Table 3. In place of directly supplying command-line arguments, it is also possible to configure shell environment variables that will automatically populate arguments in future runs of `tfmgr`. `tfmgr` will automatically upload the resources from step 1 to the OSF project from step 2. It will then add git-annex pointers to the data to the forked repository from step 3. Finally, it will open a pull request that proposes adding the new template resources to the main TEMPLATEFLOW ARCHIVE.

Supplementary Box 2. A framework where atlas knowledge flows across templates.

Researchers frequently report coordinates and define new atlases with reference to the stereotaxic coordinate spaces defined by template images. To access the information encoded in those atlases and coordinates, scientists are accordingly required to perform their own analyses with respect to the same template image. If scientists then wish to access information defined with reference to another template, they must perform further, potentially costly analytic steps.

We designed TEMPLATEFLOW to facilitate the transfer of information across the coordinate spaces defined by templates. To fulfill this end, a standardised and version-controlled archive of template datasets must be complemented with a means of mapping features between template images. These mappings take the form of bi-directional, diffeomorphic transformation functions that can be composed to project data from any template to any other template. Accordingly, TEMPLATEFLOW incorporates an image registration workflow that leverages the top-performing symmetric image normalisation (SyN) algorithm from ANTs to compute template-to-template transforms between human brain template datasets²⁴.

The TEMPLATEFLOW registration framework proceeds in several steps (Supplementary Figure 1). First, a sample of 10 high-quality⁶⁹ adult human T1-weighted images from the MNI 152 database are subject to atlas-based brain extraction (or skull-stripping), which sets all non-brain voxels to zero. The 10 skull-stripped images are then aligned via SyN to the two templates between which the transformation is being computed, thus obtaining maps of all 10 images transformed into both templates' coordinate spaces. Afterward, a single, high-quality template-to-template transform is obtained by computing a 10-channel registration to simultaneously maximise the alignment of all 10 images from one template's coordinate space into the other's.



Supplementary Figure 1. The *TEMPLATEFLOW* registration framework. To share information between human brain templates, *TEMPLATEFLOW* computes a complete set of transformations that map between all human template datasets in the *TEMPLATEFLOW* ARCHIVE. Each template-to-template transformation is computed as follows: first, 10 T1 images are skull-stripped and then registered to both templates. A single registration is then computed between the two template spaces so as to simultaneously optimise the alignment between all 10 images in both spaces.

Each template-to-template transform is a function that maps images from a source template space into a target template space. The *TEMPLATEFLOW* archival specification pairs each transform with an intuitive name that allows users to easily determine the transform's source and target. Specifically, *TEMPLATEFLOW* archives template-to-template transforms in the directories of their target templates. Transforms are stored in HDF5 format and are denoted with the suffix `xfm`, and their source templates are denoted in the `from` field of the transform file name. Transforms can be readily applied to datasets using ANTs, supporting the sharing of atlas knowledge across template datasets.

By chaining template-to-template transforms with the template-to-subject transforms typically computed in the course of a standard neuroimaging workflow, atlas knowledge can also freely flow from any *TEMPLATEFLOW* dataset into subject-specific coordinate spaces.

Supplementary Box 3. Example use case: ICA-AROMA.

Suppose that a user wishes to denoise a BOLD MR image using ICA-AROMA⁷⁰ and then perform regional analyses using atlases defined with respect to the most recent MNI152NLin2009cAsym template^{34,48}. Because ICA-AROMA operates in FSL's MNI152NLin6Asym space²³, this workflow would be nontrivial to configure, and inaccurate

configuration would incur analytic inconsistencies. TEMPLATEFlow undergirds *fMRIPrep*'s solution to this problem by providing a uniform system for retrieving resources that allow both templates to be integrated into the processing workflow⁷¹. In particular, the same code can be reused to register the functional image to both target templates. ICA decomposition and identification of motion-related components by way of the ICA-AROMA algorithm are performed in the MNI152NLin6Asym space. The time series of components identified as noise can then be filtered in the subject's native space before *fMRIPrep* projects the BOLD data into all desired target spaces, including MNI152NLin2009cAsym.

Ease of use and broad access are not without attendant challenges. The friction preventing inappropriate use of resources is arguably reduced when a wider assortment of resources is on hand, and users are afforded additional liberty and responsibility for selecting the resource most appropriate for their objective. For instance, there is no inherent mechanism preventing a researcher from using an adult template for a study of infants, although this usage would not be well-motivated. Thus, TEMPLATEFlow resources must be used with active intent and deliberation on the part of the user.