



HAL
open science

Mathematical Models for Personalized Interventional Radiology : Application to Cancer Treatment

Tom Boeken

► **To cite this version:**

Tom Boeken. Mathematical Models for Personalized Interventional Radiology : Application to Cancer Treatment. Statistics [math.ST]. Institut Polytechnique de Paris, 2024. English. NNT : 2024IP-PAX036 . tel-04735430

HAL Id: tel-04735430

<https://theses.hal.science/tel-04735430v1>

Submitted on 14 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2024IPPAX036

Thèse de doctorat



Modèles mathématiques pour la radiologie interventionnelle personnalisée : Application au traitement du cancer

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°574 de l'Institut Polytechnique de Paris, École doctorale de
mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques appliquées

Thèse présentée et soutenue à Paris, le 26 juin 2024, par

TOM BOEKEN

Composition du Jury :

Erwan Le Pennec Professeur, Ecole Polytechnique (CMAP)	Président
Joan Glaunes Maître de conférence, Université Paris Cité (MAP5)	Rapporteur
Marco Lorenzi Chargé de recherche, INRIA (EPIONE)	Rapporteur
Carole Lartizien Directrice de recherche, CNRS (CREATIS)	Examineur
Olivier Pellerin Professeur, Université Paris Cité (PARCC)	Examineur
Stéphanie Allasonnière Professeur, Université Paris Cité (HeKA)	Directeur de thèse
Marc Sapoval Professeur, Université Paris Cité (PARCC)	Co-directeur de thèse

Part One: Introduction To Our Challenges

August 14, 2024

Contents

1	Abstract	2
2	Thank you	3
3	Why Interventional Radiology	5
4	Brief history of manual interventions	6
5	How computer vision will help autonomous medicine	7
6	Challenges	8
7	Purpose of this work	8

1 Abstract

The integration of computer vision into Image-Guided interventions has the potential to change our medical practice. This work lays some bricks for the future of autonomous interventions in our specific field regarding cancer patients, addressing key components necessary for its realization. We first explore the transformative impact of AI on the physical abilities of interventional radiologists. We emphasize the need to navigate technical and ethical challenges. Interdisciplinary collaboration and robust evaluation processes are highlighted as essential for the safe integration of AI into clinical practice. We then propose an organ agnostic method for detecting focal anomalies on volumetric cross-sectional imaging. Leveraging the Large Diffeomorphic Deformation Metric Mapping (LDDMM) framework, this approach showcases enhanced object reconstruction and precise lesion localization. In the same framework, we propose a classifier, where patient selection presents unique challenges due to the complex benefice/risk ratios. To go beyond images, clinical data from tumor DNA analysis is integrated into a prospective study specifically conducted for this work. Generative Adversarial Networks (GAN) and Modelling Atlases Using the Markov Chain Monte Carlo - Stochastic Approximation Expectation-Maximization (MCMC-SAEM) Algorithms are used to predict patient trajectories. This approach enables the exploration of new trajectories, enhancing our understanding of disease progression and treatment response in relationship of circulating tumor DNA. Lastly, we explore advanced visualization techniques for in vivo and ex vivo 3D vasculature. We propose a planar representation of undescribed anatomy, offering a promising avenue for further exploration and understanding. Together, these sections offer solutions to parts of the realization of autonomous interventions in our field.

2 Thank you

This work was directed by Pr Stéphanie Allasonnière and Pr Marc Sapoval.

Thank you Stéphanie for your incredible support. Since our first *Masterclass* in 2019, you have taught me perseverance and you have opened my mind to many new horizons. I look forward to the challenges we will tackle next!

Thank you Pr Sapoval, "*patron*", for allowing me to embrace our amazing specialty beyond imagination. Thank you for making interventional radiology a passion.

Thank you Erwan Le Pennec, Joan Glaunes, Marco Lorenzi, Olivier Pellerin and Carole Lartizien for accepting the invitation to constitute the jury.

I would like to deeply thank the most perfect team, Pr Olivier Pellerin for shaping tomorrow's IO, Dr Alessandro Di Gaeta, Dr Marc Al Ahmar, Dr Xavier Guerra for setting the bar so high. Thank you Carole Déan for taking clinical research to the next level. Thank you Laetitia Thuloup, Anaëlle Creusot, Muriel Garric, Olivia Nerriec and Laura Hivert for helping us build the best department ever.

This work was made possible by Jean Feydy. Jean, your passion for utilizing mathematics to aid patients is contagious. I can't wait to implement the innovation you bring to our field into daily practice.

It was also made possible by the extended team, Louis Pujol, Alisa Kugusheva, Julien Nguyen Van, Philippine Cordelle, Vianney Debavelaere, Clément Mantoux. Thank you Sarah Zohar for leading HeKA where it all happened!

I would like to thank Pr Antoine Feydy for offering me the wisest advice. I still remember the day I walked into your office to apply for the Radiology Residency.

Thank you, Pr Olivier Clément, for providing me with the opportunity to engage in all of these projects.

Thank you, Pr Frederic Clarençon and Eimad Shotar for believing in the power of anatomy. I hope this work will provide new perspectives in your field.

Thank you, Pr Hélène Blons, for introducing me to Cancer Genetics and for your ongoing support.

Thank you, Pr Alain Chédotal, for allowing us to embark on the Embryology journey.

Without my friends and family, none of this would have any meaning. Thank you is not enough.

I dedicate this work to my fantastic wife Fanny and to our son Sacha,



Figure 1: Stéphanie visiting our unit, with Pr Pellerin.



Figure 2: The interventional team.

3 Why Interventional Radiology

This section was adapted from the following publications:

· Boeken, T., Dean, C., Pellerin, O. et al. *How Artificial Intelligence will Reshape our Interventional Units. Cardiovasc Intervent Radiol*

· Seah J, Boeken T, Sapoval M, Goh GS. *Prime Time for Artificial Intelligence in Interventional Radiology. Cardiovasc Intervent Radiol.*

· Boeken T, Feydy J, Lecler A, Soyer P, Feydy A, Barat M, Duron L. *Artificial intelligence in diagnostic and interventional radiology: Where are we now? Diagn Interv Imaging.*

· Boeken T. *Redefining challenging liver thermal ablation cases: present realities, future prospects. Clin Res Hepatol Gastroenterol.*

Computing science has already delivered on its promise to improve the quality of healthcare. Computers are now playing a crucial role in *clinical* medicine by augmenting our intellectual abilities in image-processing tasks like radiological interpretation or in extracting meaningful information from large and complex datasets, such as those generated by genome sequencing, patient data warehouses and epidemiological studies. Pitfalls aside, there is much promise.

This statement follows more than 70 years of research. Since the 1950's, computer scientists have been striving to enhance our cognitive abilities, such as visual perception and thought formation. Today, Artificial Intelligence has become an integral part of our daily lives, making it possible for everyone to experience its tangible benefits. Although computers are still far from possessing true autonomous thinking, the notion that they can simulate human-like thought processes is widely accepted.

At the forefront of attention is the potential for medicine augmented by AI to improve *clinical* capabilities. Whether it is through assisting or replacing certain tasks, AI can help enhance diagnostic accuracy (of a disease, or an epidemiological outbreak), predict treatment outcomes (through targeted therapies and patient clustering), and even facilitate clinical research. The potential offered by data-driven approaches raises societal concerns about the level of automation we desire for our healthcare system, and highlights the need to strike a balance between decision-making support and decision-making itself.

However, the impact of AI on the *physical* abilities of physicians has been relatively underexplored. This is particularly relevant in the growing field of *Interventional Radiology*, where our techniques based on dexterous hand manipulation could be revolutionized by contemporary data-driven methods. By definition, autonomous robotics implies self-decision making, and as a result, it might effectively further complexify the ongoing discussion around automated healthcare.

Automation in medicine shows great promise in delivering more precision, safety, consistency, enhanced ergonomics and even possibilities of remote interventions. With the prospect of autonomous interventions, the field is poised for a major disruption. Interventional radiologists may have to redefine the *manual* essence of their practice. In this specific field, artificial intelligence enabled solutions have both the potential to aid doctors in their existing areas of expertise and open up novel avenues for addressing unaddressed challenges.

In comparison with other interventional and procedural fields such as surgery or endoscopy, IR is data-rich. IR is one of a few specialties where a record is kept of the entire procedure in a standardized format (the DICOM format), is available retrospectively and **these datasets are mostly unexploited today.**

This is why the present work explores the possibilities of AI in our daily interventional units.

4 Brief history of manual interventions

The extent to which physicians are willing to substitute their physical abilities with AI-powered technology is uncertain. Considering the history of surgery, autonomous interventions may represent a significant paradigm shift from the patterns of technological and medical advancements that have been established for a long time. This history is marked by numerous ups-and-downs and turnovers along the way, all related to the physical ability to improve (in best cases) patient health through removing, restoring or replacing organs.

“Surgery was a desirable skill for a physician to acquire” in 13th century Southern Europe, Roy Porter argues in *The Greatest Benefit To Mankind*, “elsewhere, the gap widened and surgery was excluded from the academic curriculum”. With the creations of guilds, such as Paris’ surgeon’s organization in 1210, London’s Fellowship of Surgeons in 1368, or the Company of Barbers in 1376, surgical schools evolved along and sometimes away from clinical medicine. In particular, the Barbers had traditionally carried out minor surgical procedures, while the more educated surgeons attempted to establish surgical principles, up until 1745 when a division between the two groups occurred, leading the surgeons to form the Company of Surgeons.

In the mid-1800’s, surgery shifted from an experimental towards a rational therapeutic activity. Anesthesia and asepsis “offered the unprecedented prospect of safe and virtually unlimited surgical intervention”. In a 1933 issue of the *Annals of Surgery*, one of the founders of the American Board of Surgery writes that the period from 1880 to 1890 proved to be in surgery the concretization of the practice as a science, thanks to Lister’s “most important discovery in the development of scientific surgery, antiseptis”. The description of early 1900’s operating rooms filled with sterile instruments and surgeons covered in rubber gloves deeply resembles contemporary interventional medicine. As dexterous physicians emerged throughout the medical scene, “fame and fortune awaited the surgical pioneer who first laid the knife to some hitherto untouched part – perhaps he would be immortalized by an eponymous operation”.

During this golden era for surgery, Theodor Billroth (1826-1894), sometimes regarded as the founding father of cancer surgery and as a brilliant musician, wrote his classic textbook *General Surgical Pathology and Surgery* in the same decade he crossed paths with composer Brahms. The latter would pay tribute to the surgeon-musician by dedicating two string quartets, which Siddharta Mukherjee observed as no coincidence in *A Biography of Cancer*: “both push *manual skill* to its limit; both depend on immediacy, precision, and opposable thumbs”. It is remarkable that even with the rise of mass production and machinery away from *hand* manufacturing during the Second Industrial Revolution, manual dexterity remained the uttermost requirement for quality surgery. **Treating patients with bare hands was never perceived as anachronic.**

Technological advancements later resulted in less invasive therapies, with the first angiography of blood vessels being reported in 1896 and the first cerebral angiography in 1927. Flexible endoscopes equipped with glass-fiber optics during the same period opened the door to a new realm of interventional procedures.

In retrospect, one century later, there hasn’t been a solitary significant breakthrough that has altered so deeply the trajectory of interventional medicine since the golden 1880-1930 period. Interventional medicine, through surgery, endoscopy or interventional radiology, has revolutionized healthcare, enabling less invasive and safer procedures for the diagnosis

and treatment of various medical conditions. The integration of artificial intelligence and robotics might further propel the growth of interventional medicine, bringing about novel applications and improved patient outcomes.

5 How computer vision will help autonomous medicine

The potential clinical benefits of autonomous medicine are multiple. Improved accuracy and precision, reduced procedure times and augmented ergonomics are the key clinical findings in robot-assisted procedures. Meta-analyses are now available for each subspecialty, for example in uro-oncology, gastric surgery, rectal surgery, or thoracoscopic interventions. Though comparative data differ in terms of clinical outcomes and complication rates, most studies show a trend towards shorter hospital stays and lower conversion rates to open surgery.

Additional benefits of robotic assistance extend outside the scope of the operating room. Outer space interventions are a niche but remote surgery capabilities can expand access to expert care in underserved and rural areas, leading to better healthcare outcomes and reduced disparities. Other benefits include simulation solutions, as offered by most manufacturers, and specific training programs for physicians.

These surgery-specific robotic systems also face common limitations, including the consumption of time, resources, and space. Another challenge is the absence of tactile feedback, with human-machine interactions primarily being one-sided. Though some degree of force feedback is available, algorithms that can effectively utilize the full tactile data from existing sensors are still restricted in robotic systems. To address this limitation, multidisciplinary research teams are exploring bioinspired concepts such as dynamic bidirectional control and human-in-the-loop systems. These approaches aim to transmit high-density tactile information from a robot to an operator, replicating the natural tactile experience with enhanced fidelity.

Contemporary surgical robots are not equipped with the advanced perception and decision-making capabilities required for full autonomy. They cannot independently perceive their environment, recognize potential complications, or adapt to unforeseen situations during surgery. This level of autonomy would require significant advancements in AI, computer vision, and sensor technologies, detailed later in this work. The integration of such systems into clinical practice will need to be accompanied by careful consideration of the ethical, legal, and technical challenges involved.

The present work focuses on the pre-requisites for full automation: how to select patients prior to the intervention, and how to treat them using current devices.

There are two main pathways to use images for live guidance in interventional radiology:

- ***Endovascular approaches***: by navigating directly within the vessels, including stenting, thrombectomy, and embolization interventions.

- ***Percutaneous approaches***: by targeting lesions through the skin with live image fusion including biopsies, tumoral destruction via thermal ablation, bone cementoplasty and screw fixation.

Endovascular approaches enable physicians to navigate into the patient via the vessel network and without requiring open surgery. Live images are retrieved using fluoroscopic technology and radio-opaque devices such as catheters, guidewires, contrast media and stents. Most procedures are performed by a manual operator who inserts these devices while analyzing the live images, using a two-degree of freedom action: the push-pull action and the tork action.

Percutaneous approaches include all therapies that use a predefined trajectory to insert a needle into a target organ under live imaging. This needle can be used for diagnostic interventions (biop-

sies), or for treating tumors with thermal ablation (cryoablation, microwave or radiofrequency). The same trajectory can be used for cementoplasty or screw fixation in bone diseases.

6 Challenges

The challenges and ethical implications in autonomous interventions include ensuring the responsible development and deployment of AI technologies, addressing legal and regulatory concerns, and maintaining the vital role of human oversight in healthcare delivery. Achieving full autonomy in interventional medicine requires overcoming technical challenges, such as robust perception, decision-making, and control. Additionally, AI algorithms must be able to adapt to unforeseen situations and work effectively with human operators. As AI and robotics become more autonomous, concerns regarding liability, patient consent, and the role of human oversight must be addressed. Ensuring transparency, fairness, and robustness in AI algorithms is essential for public trust and acceptance.

The integration of AI and robotics in interventional medicine is likely to continue, with more advanced systems capable of performing complex tasks with minimal human intervention. This collaboration will require interdisciplinary research, including the development of AI algorithms, robotics hardware, and human-machine interfaces.

The evaluation of AI-integrated systems will demand rigorous preclinical testing, adapted clinical trials, regulatory approval, post-market surveillance, health technology assessment, and proper training and education for healthcare professionals. By addressing these stages and considerations, AI-integrated systems can be safely and effectively incorporated into clinical practice, ultimately improving patient care and outcomes.

7 Purpose of this work

The purpose of this work is to show how recent advancements in the Applied Mathematics field enable to build those foundations towards autonomous interventional radiology.

As an illustration, we decided to focus on colorectal cancer patients, from selection (referred from tumor boards) to live 2D and 3D-guided therapies.

The *green projects* refer to the detection and segmentation of tumors from pre-therapeutic CT-scans. Then, the *orange projects* deal with predictive biomarkers based on imaging and biological data and refer to patient selection. The *purple projects* explore the possibilities of understanding live images stemming from our interventional units. The *red projects* explore how this should help make our work more autonomous and how it will affect our basic understanding of human anatomy.

Better *patient selection* is similar to the concept of precision medicine and is at the core of interventional radiology. Decision support systems may help tailor treatment decisions based on imaging phenotypes, yielding better clinical results. Interventional radiologists rely on multidisciplinary boards for oncological treatment strategies. These board discussions perform multiparametric risk-stratification, integrating the patient's full data before a treatment is advised. Few applications replicate or even outperform these discussions by predicting the outcome from data available in each specialty (radiology, histology, molecular biology, etc.). The ability to incorporate clinical information, imaging data and genetic information should improve the objectivity and accuracy of decision-making, and, down the streamline, patient outcome. These challenges are assessed

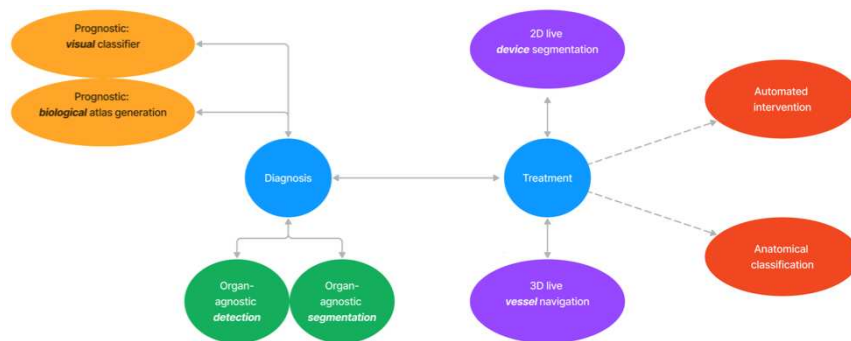


Figure 3: PhD map.

in the *orange* and *green* projects. What happens next, in our interventional unit, is where we will probably witness the most spectacular shift towards full autonomy. The *purple* projects assess how modern methods will feed future robots with thorough understanding of the live environment. *Whether we desire the future to be shaped by these advances is another issue, not explored here.*

Part Two: Organ Agnostic Detection Of Focal Anomalies On Volumetric Imaging Using The Large Deformation Diffeomorphic Metric Mapping Framework And Application Of A Classifier For Patient Selection

August 14, 2024

Contents

1	Detection of focal anomalies	3
2	Introduction	4
3	Detection of anomalies using residuals	7
3.1	LDDMM elements : Deformation model and template estimation	7
3.2	Presentation of the model	10
3.3	Computation of the template using a hypertemplate	11
3.4	Comparison to other models	11
4	Simulated example	12
4.1	Data set	12
4.2	Application of the models presented in section 3	12
4.3	On the choice to estimate anomaly matrix and deformation at the same time	14
5	Application to a data set of brains with tumors	15
5.1	Presentation of the data set	15
5.2	Results	15
5.3	Application with only one control and one sick subject	17
6	Application to the liver data set	18
6.1	Pre-processing	18
6.2	Presentation of the results	19
6.3	Quality of the detection	23
7	Limitations to the detection of focal anomalies	24
8	LDDM-based classifier from pre-therapeutic CT scans: towards patient selection	25

8.1	Why selection matters	25
8.2	Rationale for Chemoembolization	27
8.3	The FFCD-1201 trial	29
8.4	Classification based on the previous segmentation	29

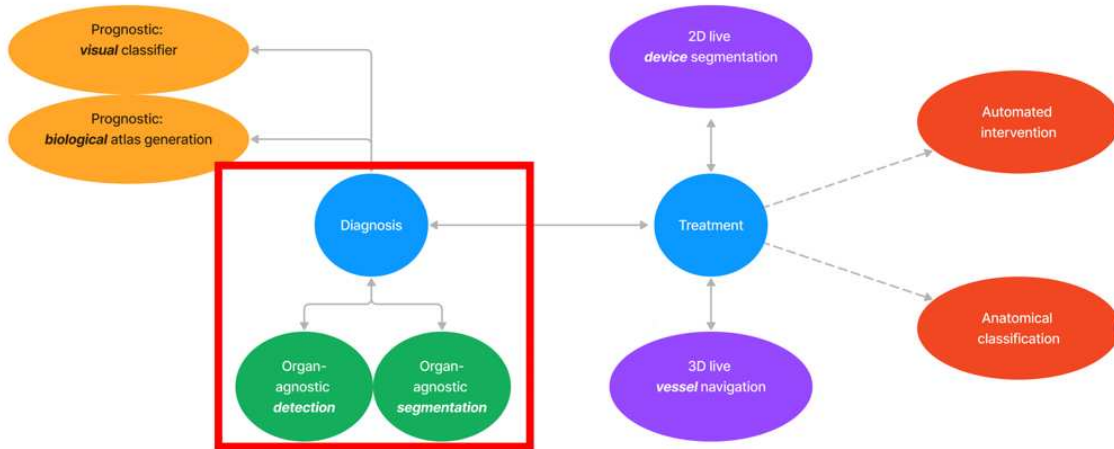


Figure 1: PhD map.

1 Detection of focal anomalies

This section refers to the *green projects*.

This section was adapted from the following publications:

- *Organ agnostic detection of focal anomalies on volumetric imaging using the large deformation diffeomorphic metric mapping framework. Submitted in IEEE Transactions on Medical Imaging*
- *Patent number 2023/247573. “Method implemented by computer means for characterizing at least one observation of a subject”*

Special thanks for this section: Vianney Debavelaere and Jean Feydy for making it accessible throughout the years.

Abstract This section introduces a method for detecting focal anomalies on volumetric cross-sectional imaging without requiring annotations or extensive training datasets. Leveraging the Large Diffeomorphic Deformation Metric Mapping (LDDMM) framework, this method is exemplified using liver CT scans and brain MRIs. The presence of focal anomalies prevents us from directly applying a deformation of a template towards a patient with an unknown number of focal anomalies. We propose to use at our advantage this diffeometric constrain by modelling the residuals (i.e. the difference between the deformed template and the observation) as a sparse matrix in addition to an independent noise. We can use the residuals of the diffeomorphic deformation from a control template to detect and segment lesions in any organ. What cannot be reconstructed as a diffeomorphic deformation of the template is hence put in this matrix and classified as an anomaly. We demonstrate that simultaneous estimation of deformations and anomalies enhances outcomes by minimizing reconstruction errors. Applied to simulated data, our method yields improved object reconstruction and accurate lesion localization. This adaptability extends to real data, where we apply the approach to liver and brain datasets. In the liver dataset, our approach detected 94 percent of metastases, showcasing its effectiveness.

Keywords Generative statistical model, Diffeomorphic deformation, sparse lesion detection, organ agnostic detection

Acknowledgment The dataset was provided by the Fédération Française de Cancérologie Digestive from the FFCD-1201 investigators group. The research leading to these results has received funding from the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

2 Introduction

This study introduces a geometrical approach for detecting and segmenting focal anomalies within organs in cross-sectional volumetric medical imaging, such as CT and MRI data. Unlike traditional methods, this approach does not require annotations or extensive training datasets. Leveraging the Large Diffeomorphic Deformation Metric Mapping (LDDMM) framework (Trouvé, 1995; Dupuis et al., 1998; Beg et al., 2005), this organ-agnostic method offers flexibility across various clinical scenarios. To showcase its effectiveness, we present applications on a liver and a brain dataset.

The study objective is to identify focal anomalies from a healthy template, using an organ agnostic approach. This differs from traditional semantic segmentation methods.

The proliferation of medical imaging and computational advancements has facilitated the creation of models for anomaly detection, classification, and segmentation. These models streamline tasks, enhance reporting accuracy, and enable quantitative disease monitoring. The accurate detection and assessment of focal anomalies in medical imaging are especially crucial due to increasing cancer rates and rising clinical workloads. Focal anomalies can be discovered incidentally or monitored quantitatively during the disease’s progression. Most models rely on labelled training data to learn to identify anomalies on images.

Deep-learning based segmentation methods contributed significantly to the specific requirements of biomedical data, especially since the emergence of U-Nets architectures. Fully convolutional neural networks are extensively evaluated for deep learning segmentation of anatomic structures, including images with focal anomalies. Recently, Wasserthal et al (Wasserthal et al. (2023)) published a model outperforming publicly available segmentation solutions, the ‘TotalSegmentator’, trained on more than a thousand labelled CT scans and tested on four thousand CT scans yielding a 0.943 Dice score. The model is based on a widely popular task-agnostic model proposed by Isensee et al (Isensee et al. (2020)), the so-called nnU-Net, ‘no new UNet’ network, a biomedical domain tool that automatically configures itself for various tasks, including preprocessing, network architecture, training, and post-processing. nnU-Net, and its comparable competitors’ performances are attributed to systematizing the manual method configuration process, without proposing a new neural network architecture.

Such methods still require large-scale datasets, and, in most cases, labelled anomalies. This limitation will commonly provoke over fitting on a low sample size training set, and failure for the model to recognize an unseen image. A recent shift towards generalizable solutions is best illustrated by the Medical Segmentation Decathlon (Antonelli et al. (2022)) which hypothesizes that a method capable of performing well on multiple tasks will generalize well to a previously unseen task. This contrasts with the previous era of custom designed models for specific segmentation tasks. It is noteworthy that more than 60 percent of the participants used a U-Net and that all methods were

based on convolutional neural networks. Building on the specific segmentation task of 3D brain tumors on MRI, recent techniques such as Hatamizadeh et al (Hatamizadeh et al. (2022)) include tuning such networks. This team proposed a Swin UNet with Transformers, where the input, comprising multiple modalities, undergoes transformation into a 1D sequence of embeddings. This sequence serves as the input for a hierarchical Swin transformer functioning as the encoder, yielding high performances.

Deep learning segmentation methods, such as Wasserthal et al.’s or Hatamizadeh et al’s, have been tailored to biomedical data, achieving performance on tasks like anatomic structure segmentation in images with focal anomalies. They emphasize on the importance of methods capable of performing well across multiple tasks without the need for task-specific model designs, or prior large labelled datasets. These approaches are bound to improve in performance with the availability of large-scale datasets specifically labelled for multi-organ segmentation such as the CT-ORG multi-organ set (Rister et al. (2020)). As demonstrated by a recently published foundation model for medical imaging based on more than 1 million patients (Ma et al. (2024)), there is a clear shift towards task-agnostic and organ-agnostic models.

Alongside the availability of large-scale datasets, the specific question of transferability – i.e. the use of one single model to perform various segmenting tasks on various imaging modalities – is being assessed with large-scale models. This is best illustrated by the works of Huang et al (Huang et al. (2023)) who propose a scalable (to different parameter sizes), transferable U-Net with parameters size ranging from several millions to one billion. This model was also pre-trained on the TotalSegmentator dataset, using the same nnU-net framework mentioned earlier.

Despite their success, U-Nets and their related solutions are not an all-encompassing solution. Modern medical shape analysis might, however, combine neural networks with “classic” atlas-based registration, paving the way towards more scalable and generalizable approaches. The convolution product is at the heart of deep learning-based shape analysis, and was naturally at the center of most studies on medical shape analysis. Such frameworks are, however, not the only high-performing approaches to organ and lesion shape analysis. Feydy et al (Feydy (2020)) provided robust implementations of fundamental geometric operations for medical shape analysis, including optimal transport, nearest-neighbor search, and image deformation. Based on the observation that CNNs remain biased towards texture, shape and pattern detection, this field proposes metric structures that are medically – and anatomically – relevant. The packages involved in this research, namely KeOps, GeomLoss, and Deformetrica, make topology-aware metrics available for processing medical shapes based on Riemannian geometry (Marsland and Sommer (2020)).

The field was structured around the idea that shape metrics could be defined in relation to registration algorithms (Durrleman et al. (2014)), best illustrated by the LDDMM framework (Schiratti et al. (2015)). Image registration can be used for the organ segmentation task, including diffeomorphic and graph-based methods. Registering consists in finding the transformation that will map a source and a reference to the same coordinate system. The development of registration methods depends on medical image data sets and frameworks specifically tuned for training and validating 3D registration methods, such as the 2023 Learn2Reg challenge (Hering et al. (2023)) (in association with MICCAI). Deep learning-based methods have emerged to address computational time limitations (Dalca et al. (2018)) of atlas-based registration, enhancing the performances of the models. Some limitations include the need for large-scale datasets to train the network and the challenge of noisy transformations. In cases involving anatomies with irregularities, such as tumor-affected areas, these techniques struggle to accurately capture volumes at specific locations due to insufficient similarity between them. This particular aspect was studied by Estienne et al

(Estienne et al. (2020)) who proposed a dual 3D CNN based architecture for joint brain registration and tumor segmentation. By computing both tasks simultaneously, they were able to use the distortion challenge posed by focal anomalies as a means for segmenting tumors.

The present study investigates the potential of segmenting focal anomalies from low sample size datasets, using a geometrical framework in an organ-agnostic detection model. This is particularly relevant for clinicians who seek to tailor the model to their specific needs when pre-trained models are unavailable.

The Large Diffeomorphic Deformation Metric Mapping (LDDMM) (Durrleman et al. (2012)) framework constructs a representative object (or template) of a population by modeling a subject as a diffeomorphic deformation of this representative object. However, the presence of tumors prevents us from directly applying it to account for deformations. Indeed, each subject can have a different number of lesions which leads to a different topology. If we directly apply the LDDMM framework, the reconstruction of each subject is obtained as a diffeomorphic deformation of the template. If this template contains dark spots, each deformation will also have dark spots, even if the targeted subject does not have this number of lesions. It may even force the model to try to make some dark spots appear or disappear by using strong deformations in areas where the anatomical shape should not change. Therefore, a model driven by a template with lesions may not be the right one. One could rather propose a typical organ for the template and lesions as additional elements not concerned by the global shape of each subject supporting the need to transform the model.

In this paper, we take advantage of this diffeomorphic constrain to automatically detect anomalies on a set of images. More generally, we suppose that we have at our disposal a data set of subjects without anomaly and another set with anomalies, each may be of low sample size. The goal is to immediately detect the presence, or not, of these anomalies.

The method we propose does not require any annotation and is organ agnostic. We suppose that we have a data set of control patients from which we are able to create a control template using the LDDMM framework. In practice, obtaining a data set of control patients is often easier, for example by considering images from patients suffering from any other pathology with no impact on the organ considered. Moreover, to estimate an atlas, the LDDMM framework does not require broad training data. Indeed, we are not estimating parameters of a blind neural network but rather parameters of a hierarchical statistical model mimicking the data generation. This control template will characterize a control population and hence, will have no anomaly. With this template fixed, anomalies on a new subject are then defined as what cannot be obtained as a diffeomorphic deformation of this control template.

To extract these anomalies, we model the residuals (i.e. the difference between the deformed template and the observation) as a sparse matrix in addition to an independent noise. What cannot be reconstructed as a diffeomorphic deformation of the template is hence put in this matrix and classified as an anomaly. The goal is to obtain these anomalies in the matrix and separate them from the noise. The idea to study the residuals is further motivated by Durrleman et al. (2011) where the authors showed that the residuals still contain information on the variability of the population. Another advantage of our model is that it retrieves the deformations from the control template towards the patients. This deformation can be interesting in the prediction of the outcome of a treatment where the environment around the anomalies sometimes plays a decisive role.

One could imagine to first estimate the deformation from the template towards the observation and only at the end of the estimation, estimate the anomalies from the residuals. However, we will show that estimating both deformations and anomalies at the same time improves the results by reducing the error of reconstruction. Indeed, considering the possibility of having additional

elements in the organ enables to see the organ globally and therefore to have a more accurate deformation.

Although the method seems to rely on a template of control patients, it can actually be estimated as well. We will propose a way to derive a template without anomaly from patients with anomalies only using a single observation of a control patient, called hypertemplate.

In the present paper, we will first present the mathematical framework, the statistical model and its estimation algorithm. Then, we apply the model on a simulated data set and show that we obtain better reconstructions of our objects and an accurate localization of the lesions. We then show its versatility by using it on two real data sets. Two distinct datasets are used to illustrate this clinical need.

The first dataset consists of CT scans collected during a Phase 2 trial involving patients with liver metastases from colorectal cancer. A considerable 80 percent of patients with colorectal cancer liver metastases have unresectable disease. To address this challenge, intrahepatic arterial delivery of chemotherapy has been suggested for patients with liver-only disease. This approach allows for concentrated delivery of anticancer agents (irinotecan) to tumor cells while minimizing systemic toxicity. In the context of this trial, intra-arterial therapies demonstrated improved treatment efficacy compared with systemic therapies. However, a third of the patients did not respond to the treatment. Detecting, segmenting, and characterizing the liver metastases could substantially enhance patient selection and enable more accurate treatment targeting. Unfortunately, performing these tasks manually is unfeasible due to their complexity. The dataset used in this study serves to exemplify the proposed method. As part of the preprocessing phase, a liver segmentation was performed, with unlabeled multiple focal anomalies corresponding to the liver metastases for all patients. The trial is registered (clinical trial number NCT01839877) and the results were published in 2020 Pernot et al. (2020).

The second dataset is derived from the publicly available Multimodal Brain Tumor Segmentation Challenge 2018 (BRATS). It includes pre-operative multimodal MRI scans of glioblastoma (GBM/HGG) and lower grade glioma (LGG), with pathologically confirmed diagnosis and available overall survival. This dataset holds importance due to its potential to aid diagnosis and survival prediction.

An example from set n°1 is given Figure 2. Tumors can be seen as slightly hypodense formations within the liver after contrast injection. The livers are segmented in the pre-processing phase but the focal anomalies are not, as shown in Figure 3 . An example from set n°2 is available here (Bakas et al., 2017, 2018; Menze et al., 2014).

As a first step, we will just consider the first observation of each patient and interest ourselves in the creation of a cross-sectional atlas.

3 Detection of anomalies using residuals

3.1 LDDMM elements : Deformation model and template estimation

Given a dataset $(y_i)_{1 \leq i \leq n}$ of images of dimension $d \in \{2, 3\}$, we want to create a template that is representative of the population. To do so, we begin by creating a distance between observations using diffeomorphic deformations.




Figure 2: Example of a contrast-enhanced abdominal CT-scan, portal phase. The red arrows indicate colorectal cancer metastases, referred as focal anomalies.




Figure 3: Segmented liver volume from initial CT scan of a patient with colorectal cancer metastases. In this pre-processing, some tumors are visible as slightly whiter volumes.

Let V be a Reproducible Kernel Hilbert Space. For $x \in \mathbb{R}^n$, a vector field v is represented as:

$$v(x) = \sum_{i=1}^{n_{cp}} K_V(c_i, x) \alpha_i \quad (1)$$

where $(c_i)_{1 \leq i \leq n_{cp}}$ are called control points and $(\alpha_i)_{1 \leq i \leq n_{cp}}$ are called momenta. v is thus represented as the interpolation of the momenta at the control points using the kernel K_V . In practice, we choose K_V to be a Gaussian kernel with variance σ_V^2 : for $x, y \in \mathbb{R}^n$,

$$K_V(x, y) = \exp\left(-\frac{\|x - y\|^2}{2\sigma_V^2}\right).$$

Let $\mathbb{L}^2([0, 1], V) = \left\{ (v_t)_{t \in [0, 1]} \mid \forall t \in [0, 1], v_t \in V \text{ and } \int_0^1 \|v_t\|_V^2 dt < \infty \right\}$.

Given $v \in \mathbb{L}^2([0, 1], V)$, we set ϕ_1^v the diffeomorphism obtained as the flow at time 1 of the

vector field v :

$$\begin{cases} \partial_t \phi_t^v = v_t \circ \phi_t^v \\ \phi_0^v = Id. \end{cases} \quad (2)$$

We then set $\mathcal{G} = \{\phi_1^v | v \in \mathbb{L}^2([0, 1], V)\}$ the group of such diffeomorphisms. It is now easy to define a distance on \mathcal{G} . For $\phi, \phi' \in \mathcal{G}$, we set:

$$d_{\mathcal{G}}(Id, \phi) = \inf \left\{ \left(\int_0^1 \|v_t\|_V^2 dt \right)^{1/2} \mid v \in \mathbb{L}^2([0, 1], V) \text{ and } \phi_1^v = \phi \right\}$$

and

$$d_{\mathcal{G}}(\phi, \phi') = d_{\mathcal{G}}(Id, \phi' \circ \phi^{-1}).$$

This exactly states that \mathcal{G} is given the structure of a manifold on which distances are computed as the length of minimal geodesic paths $(\phi_t^v)_{t \in [0, 1]}$ connecting two elements.

It has been showed that this infimum is in fact a minimum and that the distance is right invariant (Trouvé, 1995; Younes, 2010; Durrleman, 2010). This infinite dimensional problem can be approximated by a finite dimensional one where time and velocity vector field are parametrised. This construction yields a formulation of the final diffeomorphism as a function of the initial configuration of control points and related momenta (Durrleman, 2010). Therefore, a geodesic in \mathcal{G} passing through Id at the initial time is then uniquely defined by its initial velocity v_0 given by initial control points and momenta. In the following, we will write $\mathcal{E}xp_t(v_0)$ the value of this geodesic at the time t .

Hence, for two images x and $y \in M$, we set:

$$d(x, y) = \inf \left\{ \left(\int_0^1 \|v_t\|^2 dt \right)^{1/2} \mid v \in \mathbb{L}^2([0, 1], V) \text{ and } \phi_1^v \cdot x = y \right\}.$$

This distance measures the shortest length of the path relying x to y using the diffeomorphisms ϕ^v . It also allows to define a Riemannian structure on M . A geodesic on M will then be defined using an initial image p_0 and initial velocity v_0 by $t \mapsto \mathcal{E}xp_t(v_0)(p_0)$.

Hence, we measure the distance between two images as the difficulty to deform one onto another. Moreover, as ϕ_1^v is a diffeomorphism, it is invertible and preserves the smoothness and structure of the images.

We will use those geodesics to define a template of our data set, as well as the deformations from this template towards each subject using inexact matching.

In a generalization of the usual Euclidian mean, we ensure that the mean still belongs to the Riemannian space (Karcher, 1977). More precisely, we are looking for a template \bar{y} and the deformation fields v_i transforming \bar{y} onto an approximation of each observation:

$$J(\bar{y}, c_0, \alpha_i) = \frac{1}{2\sigma^2} \sum_{i=1}^n \|y_i - \mathcal{E}xp_1(v_i) \cdot \bar{y}\|_2^2 + \sum_{i=1}^n d(\bar{y}, \mathcal{E}xp_1(v_i) \cdot \bar{y}).$$

v_i is obtained using the control points c_0 and momenta α_i .

The second term $\sum_{i=1}^n d(\bar{y}, \mathcal{E}xp_1(v_i).\bar{y})$ is a regularization term while the parameter σ allows to balance between the regularity and the attachment to the data desired

\bar{y} is then the template of our population and can be estimated using usual gradient descent algorithms.

3.2 Presentation of the model

The model we propose aims at highlighting the anomalies of patients with respect to a control template. As we are modeling anomalies, we make the assumption that these features are sparse in the volume of the organ. Apart from these anomalies, the rest of the organ has to be similar to a control one. Therefore, we propose the following model.

We write $(y_i)_{1 \leq i \leq n}$ the n subjects and $\bar{y} \in \mathbb{R}^d$ the template of control subjects obtained using the methods of Section 3.1. We suppose that our images all have the same size $\prod_{j=1}^d n_j$ where $d = 2$ or 3 for $2d$ or $3d$ images.

The control points are chosen in this work along a fixed evenly spaced grid.

We assume that each observation can be written as:

$$y_i = \mathcal{E}xp_1(v_i)(\bar{y}) + A_i + \varepsilon_i,$$

where the Riemannian exponential is defined Section 3.1. This equation means that we obtain the subject y_i as the diffeomorphic deformation of the template \bar{y} : $\mathcal{E}xp_1(v_i)(\bar{y})$ to which we add a matrix A_i containing the anomalies and a noise ε_i . As explained above, v_i is a velocity field that we obtain as the interpolation of momenta $\alpha_i \in (\mathbb{R}^d)^{n_{cp}}$ at $n_{cp} \in \mathbb{N}$ control points $c_i \in (\mathbb{R}^d)^{n_{cp}}$ using a Gaussian kernel as follows: $\forall x \in \mathbb{R}^d$:

$$v_i(x) = \sum_{j=1}^{n_{cp}} \exp\left(-\frac{\|c_{i,j} - x\|^2}{2\sigma^2}\right) \alpha_{i,j}. \quad (3)$$

In order to enforce the expected sparsity of the anomalies, we add a \mathcal{L}^1 regularization on A_i . This is also a way to prevent it from including the noise, modeled as following a centered normal distribution. Here we base the regularization on Lasso regression.

Given this model, our goal is to estimate jointly the deformations from the given template and the anomaly matrices for each subject. Here, as the control template has already been estimated, we can process each subject separately and we want to minimize the following functions:

$$J_i(c_i, \alpha_i, A_i) = \frac{1}{2\sigma^2} \|y_i - \mathcal{E}xp_1(v_i)(\bar{y}) - A_i\|_2^2 + \lambda \|A_i\|_1 + \frac{1}{2} \|v_i\|_V^2, \quad (4)$$

where v_i is obtained using equation (3). The first term of J_i measures the distance between the observation and the reconstruction while the other terms measure the sparsity of A_i and the regularity of the diffeomorphic deformation.

To minimize J_i , as $\|\cdot\|_1$ is not differentiable, we implement a proximal gradient descent algorithm, using the Pytorch package to automatically compute the gradients of the differentiable part

of J_i .

3.3 Computation of the template using a hypertemplate

In some cases, it can be difficult to obtain a data set of control subjects, necessary to create the control template \bar{y} . It is for instance the case for brains where we barely get a MRI scan from a control patient. In that case, it is possible to create the template \bar{y} along the estimation of the anomaly matrices using only the image of one control patient y_0 . To do so, we consider \bar{y} to be the diffeomorphic deformation of a control hypertemplate y_0 :

$$\bar{y} = \text{Exp}_1(w_0)(y_0), \quad (5)$$

where w_0 is obtained as the interpolation of momenta β_0 at control points c_0 , as explained above. As y_0 is a control subject, it has no anomaly, and its diffeomorphic deformation \bar{y} will have no anomaly either.

Hence, in that case, we will not only estimate the velocities $(w_i)_{1 \leq i \leq n}$ and sparse matrices $(A_i)_{1 \leq i \leq n}$ but also the velocity w_0 by minimizing:

$$\begin{aligned} \tilde{J}(c_0, (\beta_i)_{0 \leq i \leq n}, (A_i)_{1 \leq i \leq n}) = & \sum_{i=1}^n \left(\frac{1}{2\sigma^2} \|y_i - \text{Exp}_1(w_i)(\bar{y}) - A_i\|_2^2 + \lambda \|A_i\|_1 + \frac{1}{2} \|w_i\|_V^2 \right) \\ & + \frac{1}{2} \|w_0\|_V^2, \end{aligned} \quad (6)$$

where \bar{y} is obtained using the velocity w_0 and w_i is obtained as the interpolation of the momenta β_i at the control points c_0 .

Note that, as the template is a reference image for the whole population, the energy to minimize is not separable and involves the contributions of all subjects.

Once again, this optimization is done by proximal gradient descent.

Remark 3.1. *Here, we use the same control points c_0 for all velocity fields $(w_i)_{0 \leq i \leq n}$ in order to reduce the computation time.*

Remark 3.2. *In practice, we sometimes have a computational problem to estimate the parameters of those models. Indeed, the gradient descent tends to stay blocked in local minima. By including the errors of reconstruction in the anomaly matrix, the algorithm often chooses not to improve the reconstruction and stays blocked in a local minimum. To solve this problem, we prevent the anomaly matrix to take any value outside of the reconstruction of the object for the first 100 iterations. This allows to improve the diffeomorphic reconstruction and to reach a more relevant area of interest of the energy landscape.*

3.4 Comparison to other models

In the following, we will compare this model to the usual cross-sectional atlas, estimating the template \bar{y} and the deformations towards the subjects without the use of anomaly matrices nor hypertemplate. This alternative method writes as minimizing the following energy:

$$F_0(\bar{y}, c_0, (\alpha_i)_{1 \leq i \leq n}) = \frac{1}{2\sigma^2} \sum_{i=1}^n \|y_i - \text{Exp}_1(v_i)(\bar{y})\|_2^2 + \frac{1}{2} \sum_{i=1}^n \|v_i\|_V^2. \quad (7)$$

We will also compare it to the cross-sectional atlas when the template is obtained via a control hypertemplate. As above, v_0 will be obtained as the interpolation of momenta α_0 at control points c_0 . The functional to minimize is now:

$$F_1(c_0, (\alpha_i)_{0 \leq i \leq n}) = \frac{1}{2\sigma^2} \sum_{i=1}^n \|y_i - \mathcal{E}xp_1(v_i)(\bar{y})\|_2^2 + \frac{1}{2} \sum_{i=0}^n \|v_i\|_V^2, \quad (8)$$

where \bar{y} is defined by equation (5).

The estimation of those two models can be done using a usual gradient descent.

4 Simulated example

4.1 Data set

To test the model, we create a simulated data set of 500 subjects deformed from a common template to which we add a random number of dark spots (between 1 and 5) and some Gaussian noise. This template is created as the deformation of an ellipse (hypertemplate). We also create a "control" data set of 100 subjects without dark spot from the same template to be able to estimate a control template \bar{y} and use it in the case of the model (4). The template and five subjects with dark spots are presented on the first line of Figure 4.

4.2 Application of the models presented in section 3

We first apply the model where the template is directly estimated, without the use of a hypertemplate nor sparse matrices (Equation (7)). As can be seen on the second row of Figure 4, a dark shadow is created on the estimated template. Similarly, this dark shadow is reported on the reconstructions of each subject. Moreover, as can be seen on the last two columns, to minimize J_0 the algorithm sometimes badly estimates the deformed object in order not to include a dark spot.

On the third row of Figure 4, we apply the model where the template is obtained from an ellipse hypertemplate but without any anomaly matrix (model (8)). This time, as expected, there is no dark shadow on the template but the reconstruction of the last two subjects is still bad. The model would rather erase one part of the shape so that the dark anomalies close to the border are considered as background. Not only removing the inside anomalies, the overall shape of the organ is badly reconstructed.

We also test the model (4). To do so, we begin by estimating a template from the 100 "control" subjects. We then fix this template in the minimization of Equation (4). This estimated template and some reconstructions are represented on the fourth line of figure 4. This time, the subjects are better reconstructed and the dark spots retrieved. Their intensity is a bit weaker than initially due to the proximal gradient descent applying a soft threshold on the residuals. However, the shapes of the dark spots are well identified and so are their positions.

Then, we apply the model (6) where the template is estimated from a hypertemplate. The results are presented on the last row of Figure 4. Once again, the subjects are well reconstructed

and the dark spots are retrieved in the anomaly matrix. Shapes, positions and volumes of the dark spots are captured.

Finally, we compare the four models by computing the mean error of registration when one only considers the form (i.e. the error between the observations without their dark spots and the reconstructions without their anomaly matrix) on Table 1. As expected, the error is smaller when considering the models (4) and (6). One can notice that the previously estimated template enables to achieve better results. This was expected as the template is better representing the control population when estimated as a representative group of images. However, when not available, the combined model reaches very interesting performances. This enhances the strength of this estimation process both in terms of template shape and anomaly detection.

	Model (7)	Model (8)	Model (4)	Model (6)
Error of registration	19,8% \pm 0.06	16.5% \pm 0.10	11.9% \pm 0.07	12.9% \pm 0.03

Table 1: Mean and standard deviation of the error of registration when one does not consider the dark spots.




Figure 4: On the top line, the template and five observations. On the second line, the estimated template and reconstructed subjects without the use of hypertemplate nor anomaly matrix (model (7)). On the third line, the results when one uses the hypertemplate but no anomaly matrix (model (8)). On the fourth line, we first estimate a template (first column) from control subjects and then reconstruct the other subjects using an anomaly matrix (model (4)). On the last line, the results with sparse matrices when the estimation of a template is done at the same time, from a hypertemplate (model (6))

Hence, we have been able to reconstruct the subjects and to retrieve their anomalies using the models (4) and (6). Moreover, we have been able to highlight an improvement of reconstruction when we estimate both anomaly matrix and deformations.

4.3 On the choice to estimate anomaly matrix and deformation at the same time

The choice to estimate both the deformation v_i and the anomaly matrix A_i at the same time, and not one after the other, comes from an effort to improve the reconstruction of the object. Indeed, we can see on the Table 1 that the errors of reconstruction are smaller when we estimate both at the same time. It can particularly be seen on the third line of Figure 4 where, for the last two columns, the residuals contain whole parts of the observations. Those parts would hence be retrieved in the anomaly matrix.

We present a last example to emphasize this need to estimate both the deformation and the anomaly matrix at the same time. The different images of that example can be seen Figure 5. This time, we add a little black line on the control template, mimicking for example a vein in a liver or a gyrus in a brain slice. From this template we create one subject to which we add a dark spot. We try to reconstruct this subject from the template, either with or without anomaly matrix. Without anomaly matrix, the model chooses to heavily deform the black line to create the dark spot. In particular, a part of the line would here be in the residual and the black spot would not entirely be in it. The estimation of an anomaly matrix from this residual would hence be bad. However, if we choose to estimate both deformation and anomaly matrix at the same time, the black line is well registered from the template to the individual and only the black spot is retrieved in the anomaly matrix.

Those two observations confirm the need to couple the estimation of anomaly matrices and deformations.




Figure 5: From left to right, the fixed template, the subject, the reconstruction without anomaly matrix and the reconstruction with anomaly matrix. Without anomaly matrix, the algorithm uses the black line to recreate the dark spot, creating a fuzzy black zone.

5 Application to a data set of brains with tumors

5.1 Presentation of the data set

We choose to first apply our model to a data set of brains with tumors obtained from the open access BraTS 2018 data set (Bakas et al., 2017, 2018; Menze et al., 2014). More precisely, the data set is composed of 50 post-contrast T1-weighted MRI scans of glioblastoma and lower grade glioma. We do not dispose of a data set of control subjects but only of the observation of one control subject. We will hence use the model (6) to estimate the template using this control subject as hypertemplate.

The goal is to reconstruct the brain of each subject from a control template and to obtain the tumors in the anomaly matrix. Ideally, the diffeomorphic deformation will register the brain folds (called gyri) and ventricles. What cannot be retrieved in the diffeomorphic deformation should hence only be the tumors.

5.2 Results

We here present the results of our model applied to this data set. On Figure 6, we present the template estimated by our algorithm.




Figure 6: The template estimated using a control patient as a hypertemplate.

On Figure 7, we show the results for four subjects. As can be seen, the lesions are retrieved in the anomaly matrix with only small errors of reconstruction. In particular, we can see that the gyri have been well registered and are not present in the anomaly matrix. As for the ventricles, if a part of one is present in the bottom right image, they are also well registered for the other images. In fact, for the bottom right patient, its right ventricle is quite different from the template, causing our algorithm difficulties to register a part of it. But, in all cases, the use of the LDDMM framework has allowed to register most of the parts of the brain without anomaly and so to obtain




Figure 7: The results for four different subjects. Each time, we put on the left, the observation and on the right, the anomaly matrix estimated. Each time, the lesions are well retrieved.




Figure 8: The patient and the anomaly matrix in a case where the tumor (red zone) is not detected.

a clean anomaly matrix.

Note that the energies we optimize are parametric. This means that we need to calibrate both σ^2 and λ . σ^2 is equivalent to the voxel noise variance if we transpose our energy to a log likelihood. As for λ , it balances the weight of the sparsity matrix with reference to the data term penalized by the smoothness of the diffeomorphism. Here, one important choice of parameter in Equation (6) is in fact the sparsity constant λ . One could choose to take a smaller λ . This would allow to include the peritumoral edema (dark area around the tumor) in the anomaly matrix. However, we would then also include more reconstruction errors. Here, we have chosen λ in order to visually detect the tumors but with as little reconstruction errors as possible, even if the whole tumor is not in the anomaly matrix. It fulfils our goal as we wanted to be able to inform the doctors of possible anomalies, which are indeed included in the anomaly matrix. We prefer to have less false positive but with more accurate location.

For a quantitative evaluation we compute the detection ratio of the 62 tumors in the data set, and we find that 59 are visible in the anomaly matrix (95%). As for the tumors not visible, they are small lesions in a zone of high variability of the brain. We show such a case in Figure 8 where the tumor is less easy to distinguish and in the middle of the brain gyri.

5.3 Application with only one control and one sick subject

Our approach here is particularly interesting as it does not require a large data set to be trained on and there is no annotation on the position of the tumors required. To highlight this advantage, we apply the exact same algorithm to only one brain with tumors. As for the template, it is fixed as a brain without tumor. Hence, we only use two different brains to try to detect an anomaly. In particular, we compare the results with the anomaly matrix estimated for this patient in the previous section where a template was estimated alongside (see Figure 9).

The tumor is once again retrieved in the anomaly matrix with small errors of reconstruction, particularly on the border and top of the brain. The errors on the top, in particular, are not present when one estimates a template alongside the anomaly matrix. In fact, the variability between the control subject and the one with tumor is bigger there, causing bigger errors of reconstruction. But, estimating a template, even with few subjects as done section 5.2, prevents this issue.

Even if more errors are included in the sparse matrix, it must be emphasized that it has been produced using only two subjects and that it shows that our method can yield usable results even without access to more than a few subjects. It also shows that a relevant pre-estimated template of control organs enables to refine the results as it takes into account a wider variability and is more representative of a control population.




Figure 9: On the left, the observation. In the center, the anomaly matrix estimated with the template fixed as a control subject. On the right, when one estimates the template, as done section 5.2. The results are showed for two different slices: at the position of the tumor (top images) and at the top of the brain (bottom images). In both cases, the tumor is retrieved. More reconstruction errors are included when one does not estimate a template.

6 Application to the liver data set

6.1 Pre-processing

We now apply our algorithm to the liver data set provided by the Fédération Française de Cancérologie Digestive regarding the FFCD-1201 registered clinical trial number NCT01839877 with specific data use agreement for the present work. Note that the full livers were manually segmented for the trial. The goal is to reconstruct the liver of each patient while recovering the tumors in the anomaly matrix. As we do not have access to a full data set of control patients, we cannot use the model (4) and we need to estimate the template from a hypertemplate using model (6). The hypertemplate is chosen as a patient, not in the database, and without tumor.




Figure 10: Example of a segmented liver.

Two different structures to appear on the livers were tumors as dark spots and vessels as white structures because tumors have lower density than normal liver parenchyma and vessels are enhanced by contrast medium. We expect to recover these separate structures in the anomaly matrix. If one only wants to retrieve the tumors, it is easy to separate them from the vessels according to their intensity. Moreover, from one subject to the other, the noise level can be totally different. If we do not preprocess the data set, we would not be able to find a sparsity constant λ efficient for each subject and, for those with the highest level of noise, this noise would be recovered in the anomaly matrix. To prevent this phenomenon, we decide to first convolve our observations with a Gaussian kernel. The resulting images can be seen Figure 11. This smoothing allows to have a robust algorithm for this population.

Moreover, because all images do not have the same pixel spacing, we need to down-sample some of the images. We also include all of them in a black box of the same size. In the following, we present the results for different subjects.

We also choose, as post processing, to put the coefficients of the anomaly matrix to 0 outside of the reconstructions and target shapes. If one does not make this choice, the errors of reconstruction are reported in the anomaly matrix. In particular, one would find white zones in the anomaly matrix at voxels where the diffeomorphic deformation has not been able to recreate a liver part and black zones where the diffeomorphic deformation has created a liver part at a place there should not be one.

6.2 Presentation of the results

We begin by presenting the estimated template \bar{y} in figure 13. Note that this template appears smoothed. The reason comes from the low number of patients with reference to the high variability in shape of this organ. Moreover, because it is derived from a particular control subject, the ves-




Figure 11: On the left, the initial images. On the right, the same subjects after convolution with a Gaussian kernel.

sels of this particular subject are still present in the final template. This may influence the future estimation of the anomaly matrix by either using these structures to force the deformations or by using the sparse matrix to erase them. A post-control should be done with this in mind. Having a template of control subjects would reduce the errors of reconstruction and allow a better detection of the anomalies.

Then, on Figure 14, we show the results for a subject without any vessel visible on the scanner. The algorithm is able to retrieve the tumors in the anomaly matrix. We can remark that the outline of the liver is also present in the anomaly matrix. It is in fact used to correct the small errors of the diffeomorphic deformation from the template towards the image and obtain a better final reconstruction. It is however easy to filter that outline to only keep the tumors using its intensity.

We then present a subject with vessels visible in Figure 15. As can be seen, not only the tumors are retrieved but also the vessels. If one wants to only find the tumors, a first possibility is to look at the negative values of the anomaly matrix. Further investigation on a post process would be




Figure 12: Control patient.




Figure 13: Template estimated as a diffeomorphic deformation of a control patient.




Figure 14: On the left, the observation. In the center, the reconstruction. On the right, the anomaly matrix.

required to only retrieve the anomalies without the small errors of reconstruction.

Finally, we show the importance to apply a Gaussian convolution to the data set before esti-




Figure 15: From left to right, the observation, the anomaly matrix, and the negative values of the anomaly matrix.

imating the parameters of the model figure 16. If one does not perform this preprocessing, the noise is retrieved in the anomaly matrix and the tumors are not visible. This problem is indeed solved after convolution and the tumor (at the top left of the liver) is retrieved.

6.3 Quality of the detection

To measure the quality of the detection, we asked a MD Radiologist to segment the tumors of 10 patients. This led to a total number of 133 tumors segmented. As we focus only on tumor lesions, we will only look at the negative coefficients of the sparse matrices to measure the quality of detection as tumors are dark spots on the liver.

We choose not to evaluate the segmentation but the detection. In fact, here, the dice score would be average as our algorithm rarely segments the whole tumor. However, as written previously, we do not want to segment the exact tumor but only to inform the medical doctor of a possible anomaly and particularly detect the location of very small lesions.

On the 133 tumors segmented, our algorithm detects 125 of them (94%). The detection was defined a non-empty clusters from the anomaly matrix above a predefined threshold size for considering an anomaly. As for the tumors which are not detected, there are two possibilities. Sometimes, the difference between the tumor intensity and the noise is really small and the algorithm is not able to separate them, in particular for some subjects for which the noise is still high. The other possibility is when the diffeomorphic registration of the liver is not perfect and the tumor is outside of it. In that case, the tumor will in fact be in the anomaly matrix but it will be lost in the error of reconstruction.




Figure 16: On the top line, the results if we do not convolve the data set with a Gaussian kernel beforehand. On the bottom, with preprocessing. On the left, the observation. In the center, the reconstruction. On the right, the anomaly matrix.

Finally, not only the tumors are retrieved in the anomaly matrix. As showed above, small errors of reconstruction can be present on the boundary. But our algorithm also plays its part by detecting other anomalies than lesions. In particular, on several subjects, some slightly dark spots are retrieved and are in fact due to a perfusion disorder.

7 Limitations to the detection of focal anomalies

In this section, we have demonstrated the efficacy of utilizing the residuals derived from diffeomorphic deformation from a control template for the purpose of detecting and segmenting lesions within a given organ. Notably, our findings indicate that this approach not only facilitates lesion identification but also enhances the overall diffeomorphic reconstruction of observed data. An inherent advantage of our method lies in its independence from extensive datasets of patient information

or annotations from medical professionals, making it particularly well-suited for scenarios where acquiring large datasets is challenging, such as in specific treatment protocols.

Our experimentation involved assessing the performance of this methodology on datasets featuring brains afflicted with glioma and livers hosting colorectal cancer metastases. Notably, in the case of glioma-afflicted brains, our algorithm successfully registered gyri and ventricles while accurately identifying tumors within the anomaly matrix.

However, several challenges persist that warrant further attention. Firstly, while our method effectively detects tumors, it also captures minor reconstruction errors. To address this, post-processing steps are necessary to isolate and extract only the relevant lesions. An initial consideration involves employing opening operations, although this approach carries the potential risk of inadvertently discarding smaller anomalies during post-processing. Additionally, a more refined strategy for handling reconstruction errors is imperative, as the absence of a tumor in the reconstructed object may lead to its inadvertent exclusion. Further refinement through post-processing of the anomaly matrix is essential to recover any overlooked anomalies and ensure the precision of our detection method.

8 LDDM-based classifier from pre-therapeutic CT scans: towards patient selection

This section refers to the *orange projects*. This section was adapted from the following publications:

Pellerin O, Boeken T, Guiu B. Chemoembolization of HCC: Time for Technical Standardization, or Is It Too Late? *Cardiovasc Intervent Radiol*. 2019 Pellerin et al. (2019)

Boeken T, Gallois C, Douard R, Méatchi T, Martelli N, Sapoval M, Taieb J, Pellerin O. Bridge to surgery after irinotecan-based liver chemoembolization for metastatic gastric adenocarcinoma: Letter to the editor. *Clin Res Hepatol Gastroenterol*. 2021 Boeken et al. (2021)

Kedra A, Boeken T, Di Gaeta A, Querub C, Al Ahmar M, Déan C, Sapoval M, Pellerin O. Exploring a Novel Technique to Tackle the Shortage of Devices for Hepatic Arterial Infusion Chemotherapy: Early Results of an Alternate Approach for Percutaneous Arterial Port Catheter Placement. *Cancers (Basel)*. 2023 Kedra et al. (2023)

Boeken T, Moussa N, Pernot S, Abed A, Dean C, Taieb J, Sapoval M, Pellerin O. Does Bead Size Affect Patient Outcome in Irinotecan-Loaded Beads Chemoembolization Plus Systemic Chemotherapy Regimens for Liver-Dominant Colorectal Cancer? Results of an Observational Study. *Cardiovasc Intervent Radiol*. 2020 Boeken et al. (2020)

Boeken T, Sapoval M, Pellerin O. Stressing the need for prolonged immediate surveillance after chemoembolization for liver metastases using drug-eluted beads with irinotecan. *Diagn Interv Imaging*. 2020 Boeken et al. (2020)

8.1 Why selection matters

Colorectal cancer (CRC) ranks as the third most diagnosed cancer globally, with its incidence projected to increase by 60% by 2030. The rapid rise in cases is attributed, at least in part, to dietary changes and the increase in obesity. A causal relationship has been suggested by a 2017 analysis involving 68 million individuals.




Figure 17: PhD map.

Metastatic colorectal cancer affects 40-60% of patients. Specifically, about 20-25% of patients present with liver metastases during their initial assessment. The management of metastatic colorectal cancer to the liver (mhCRC) is a significant challenge, closely dependent on whether the liver is the dominant site of involvement (liver-dominant disease) or if the disease is limited extra-hepatically.

Despite the prevalence of mhCRC, consensus guidelines do not agree on a unified or standardized approach. The role of intra-arterial therapies, such as those offered by interventional radiology, remains unclear.

According to the 2020 French consensus, the management starts with clinical evaluation (World Health Organization status, patient's preference), morphological assessment (CT scan and/or MRI as needed), and determination of tumor Kirsten rat sarcoma virus (KRAS), B-Raf sérine/thréonine kinase (BRAF), and microsatellite instability (MSI) status. The Thesaurus emphasizes that only resection (surgical or local ablation) can lead to complete remission. In addition to surgical and ablative treatments, the current management of patients with mhCRC includes systemic chemotherapy (5-FU, oxaliplatin, irinotecan), targeted therapies (Vascular and epithelial growth factors inhibitors, anti-VEGF, anti-EGFR), and immunotherapy (notably programmed cell death 1 or anti PD1/PDL1 based on MSI status). As third-line options, trifluridine/tipiracil and regorafenib are also available. More recently, tumor-specific characteristics have paved the way for the use of new targeted therapies. These are currently under evaluation, including anti-BRAF (encorafenib), anti-HER2 (human epidermal growth factor 2 and trastuzumab), anti-tropomyosin (NTRK) (larotrectinib, entrectinib), and anti-KRAS (sotorasib).

These therapies have improved patient management, prolonging overall survival and offering the possibility of secondary resection; this could be achieved in up to 50% of cases.

In addition, there are intra-arterial therapies proposed by interventional radiology (chemoembolization, radioembolization, hepatic intra-arterial catheter). These are included in European recommendations (ESMO) as "ablative" therapies.

Intra-arterial therapies are mentioned as Grade B or C options in the French recommendations

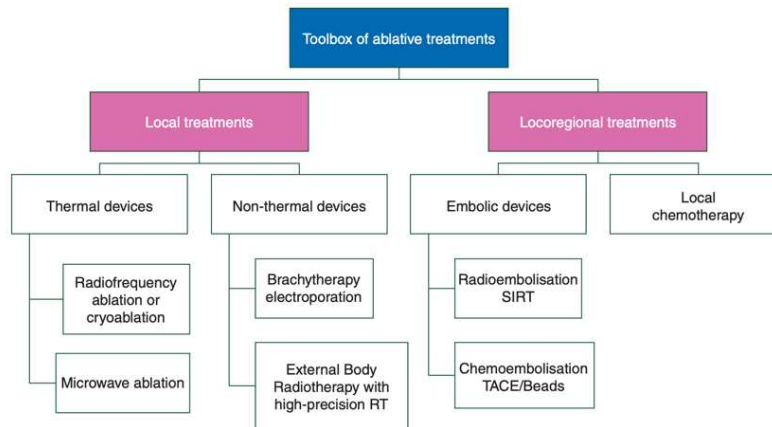


Figure 18: ESMO guidelines for CRC. High presence for IR therapies

without specific details. These recommendations state: "Whenever there are predominant hepatic metastases from CRC that are not resectable initially, the question of intra-arterial treatment must be systematically considered (approximately 15% of patients)." An international meta-analysis of 23 prospective studies highlights the challenge in establishing a unified strategy for intra-arterial treatments of colorectal cancer.

Better selection means optimizing the current place for chemoembolization in the arsenal. The approach studied here, the so-called 'DEBIRI' treatment, is later detailed specifically.

8.2 Rationale for Chemoembolization

Hepatic chemoembolization's fundamental principle rests on the liver's dual vascularization – primarily portal for healthy parenchyma and arterial for tumor lesions. As described by Bierman et al. in 1951, liver neoplasms exhibit hyperarterialization, while non-tumor liver maintains a portal-to-arterial network ratio close to 80%/20%. These findings still hold true after 70 years. The first ligation of the hepatic artery to induce tumor necrosis was reported in 1966, and the first chemoembolization likely occurred in 1974.

Chemoembolization addresses a triple challenge:

- i) Achieving maximal intratumoral chemotherapy concentration
- ii) Avoiding systemic toxicity peaks
- iii) Inducing local embolization with necrotic intent.

This procedure typically employs an emulsion of chemotherapy and lipiodol, a synthetic oil developed in 1901, with a ratio determining embolic potential and local treatment distribution.

The DEBIRI technique is an alternative to lipiodol-based chemoembolization, using drug-eluting microspheres. These microspheres act as carriers, loaded during preparation and subsequently releasing treatment closer to the target. In vitro and in vivo studies have detailed their properties and action mechanism.

In vitro physicochemical analyses of doxorubicin- or irinotecan-loaded microspheres show that these microspheres enable complete and rapid chemotherapy loading during preparation. Subse-




Figure 19: Rationale for DEBIRI, explained by Lewis et al

quent release is sustained and influenced by local ionic balance. The rationale for in vivo use of irinotecan-loaded microspheres (DEBIRI) builds upon the selective and specific tumor targeting achievable via intrahepatic arterial injection, which spares non-tumor liver primarily supplied by the portal system. Besides serving as chemotherapy carriers, these microspheres induce embolization and necrosis in tissues predominantly vascularized by arteries.

The interest of chemotherapy in the DEBIRI process has been extensively studied due to irinotecan's status as a prodrug, theoretically necessitating metabolism by healthy hepatocytes to exert its antitumor effects.

The theory is supported by two observations:

i) Concerning healthy hepatocytes: in essence, irinotecan is hydrolyzed into SN-38 by CES-2, a protein present in hepatocytes. SN-38 exhibits a 1000-fold stronger affinity for TOPO-1, the final target of chemotherapy and a crucial element in DNA replication. Unlike conventional chemoembolization techniques, DEBIRI is administered lobularly to ensure these chemical reactions by healthy hepatocytes.

ii) The second observation focuses on tumor cell pH: microspheres induce tumor hypoxia, leading to acidosis that enhances SN-38 activity.

An animal study conducted in 2011 with 54 VX2 rabbits compared the pharmacokinetics of intravenous irinotecan, intrahepatic arterial irinotecan, and doxorubicin-eluting microspheres (DEBIRI). Intra-tumoral irinotecan concentrations at 24 hours were significantly higher with microspheres versus free injection. However, these results should be interpreted cautiously since VX2 rabbits' CES-2 doesn't exhibit the same enzymatic kinetics as in humans. Similar studies were conducted for human pharmacokinetics. Intrahepatic arterial injection seems to increase systemic concentrations of active irinotecan metabolites, although publications disagree on this point. CES present in healthy hepatocytes likely explains this phenomenon. The intra-tumoral concentration of metabolites, in the absence of CES, remains less studied.




Figure 20: FFCD 1201 trial survival curve

8.3 The FFCD-1201 trial

Clinical studies on DEBIRI have evaluated treatment efficacy in palliative settings as well as in earlier stages, as first-line or neoadjuvant therapy. Most studies are retrospective. Three randomized trials were published in 2012, 2015, and 2019.

The most recent prospective study is the French FFCD 1201 phase II trial. **The investigators have granted us the access to the full data according to French law.**

This study examined DEBIRI 100-300 um (+ chemotherapy) as first-line therapy for patients with unresectable disease. The 9-month progression-free survival (PFS) was 53.6% (95% CI, 41.8-65.1%), objective response according to RECIST 1.1 was 73.2%, and R0 secondary resection was 33%. Median overall survival was 37.4 months (95% CI, 25.7-45.8), and median PFS was 10.8 months (95% CI, 8.2-12.3). In other words, one third of the patients did not respond to this treatment, while two thirds responded.

The 9-months progression-free survival was 54%: we used this to label retrospectively patients as good-responders and poor-responders in order to try and select those who would benefit most from the technique.

8.4 Classification based on the previous segmentation

Tumors were extracted using the previous organ agnostic model. The data was labelled according to clinical outcomes at 9 months as good responders (objective response) and poor responders (progression).

Good responders were also more precisely defined by the type of response and labelled accordingly. The dataset consists of the 54 extracted tumors as mesh files, and the corresponding labels. 3 patients were dropped because of unavailable data.

The LDDMM framework can be used as a classifier. A quick example is provided here. Given a set of hand-written digits ‘2’ (set A) and hand-written digits ‘3’ (set B), we re-create an atlas for 2 and an atlas for 3. A subset of the MNIST handwritten digits dataset was used.

The unknown class used here is a ‘2’ digit from outside the initial set (the subject 0). By




Figure 21: Example of a **poor responder** (corresponding label: 9-months PFS). Segmented liver and automatic detection of corresponding tumors.




Figure 22: Example of a **good responder** (corresponding label: 9-months PFS). Segmented liver, automatic detection of tumors and cartography.




Figure 23: Distribution of responses in the FFCD-1201 trial.




Figure 24: Digits set A and B from the MNIST dataset.



Figure 25: Template deformation towards unknown subject.

shooting the template towards this subject, we define its registration and related distance to the atlas. For this, we use an ambient-space dimension of 2, a kernel-width of 2, and a noise-std of 0.1, with a deterministic atlas model as presented before. Bayesian models can also be used (where, in addition to the template and the registrations, the variability of the geometry and the data noise are learned).

The estimation method used was composed of the loss function from the L-BFGS algorithm with 100 iterations. The L-BFGS estimator converges faster, and is an alternative to the Gradient Ascent estimator that might be useful in specific situations. The residuals are used here to evaluate the distance between the targets and the reconstructed objects.

We can create a classifier based on the residuals and momenta vectors of each object from set A and set B with a fixed template. We can estimate the probability of the template being closer to set A than set B with a T Student test if the sets are big enough. The examples below show an accurate classification of a ‘2’ digit and of a ‘3’ digit with a high probability. The residuals can be interpreted as the difficulty to reconstruct the deformation from the template (non-deformable portion of the template) and the momenta norms as the deformation costs.

By applying this to the liver sets, we can first compute the model with a fixed normal liver template and estimate the registration (and distance) of each patient to this tumor-free liver. We first pre-process the livers to create a surface mesh of 10.000 cells. We use an ambient-space dimension of 3, a kernel-width of 20, and a sparse-constraint of 0.1, with a deterministic atlas model with anomaly as previously described and a Proximal Gradient Ascent estimator for 1000 iterations.

The distribution of residuals to the template is provided below. The Kolmogorov-Smirnov test did not provide conclusive results towards a class but the ANOVA test did under the assumption that the distribution is normal (not confirmed here with the Shapiro test). The template used here was not a pivotal liver that enabled a classification by residuals and vector norms using this method.




Figure 26: Residuals from subject '2' to Set A and subject '3' to Set B. Note that $p\text{-value} < 0.001$.




Figure 27: Reconstruction of a '2' observation from '3' template.




Figure 28: Momenta norms of the deformations from template to set A and set B. Note that $p\text{-value} = 0.0002$.




Figure 29: Preprocessed 10.000 cells mesh of a real liver with tumors.




Figure 30: Distribution of momenta norms of all patients with a real fixed normal liver template.




Figure 31: Distribution by clinical response class from real data.

We can also use the subset of 50 simulated livers and *randomly* assign an overall response label to each patient. We compute the model with a kernel-width of 10, and a deterministic atlas model *with anomaly* as previously described and a Gradient Ascent estimator for 100 iterations. The consistency of the transformations can be found in the results: deformations are equally distant to both sets because of the random assignment of the labels concerning overall survival (p-value = 0.597 after Student-s T-Test). The LDDMM framework – and every other model – can only generate a classifier if the information is in the image provided by the pre-interventional CT assessment. It will not produce a classifier if the success rate is independent from this image, or randomly distributed.

In conclusion, we cannot conclude that the initial CT provides enough information to determine whether the patient might benefit from our technique. We did not succeed in building a classifier based on the shapes of the liver metastases. Further research should probably integrate peri-tumoral normal liver parenchyma given the high importance of normal liver for the biological mechanisms underlying irinotecan-based chemo embolization.




Figure 32: Set A from simulated set.

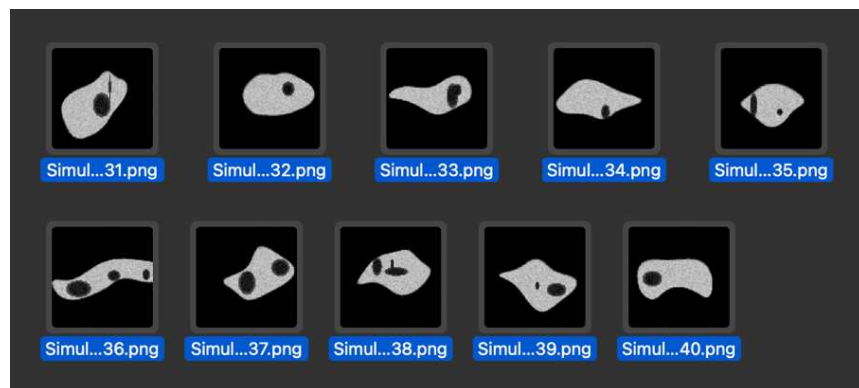


Figure 33: Set B from simulated set.

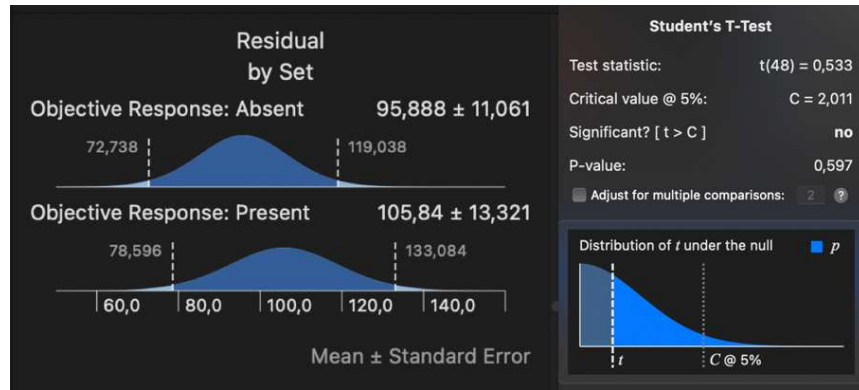


Figure 34: Distribution of residuals according to clinical response in the simulated random set.




Figure 35: Distribution of momenta norms according to clinical response in the simulated random set.

References

- Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A. Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M. Summers, Bram van Ginneken, Michel Bilello, Patrick Bilic, Patrick F. Christ, Richard K. G. Do, Marc J. Gollub, Stephan H. Heckers, Henkjan Huisman, William R. Jarnagin, Maureen K. McHugo, Sandy Napel, Jennifer S. Golia Pernicka, Kawal Rhode, Catalina Tobon-Gomez, Eugene Vorontsov, James A. Meakin, Sebastien Ourselin, Manuel Wiesenfath, Pablo Arbeláez, Byeonguk Bae, Sihong Chen, Laura Daza, Jianjiang Feng, Baochun He, Fabian Isensee, Yuanfeng Ji, Fucang Jia, Ildoo Kim, Klaus Maier-Hein, Dorit Merhof, Akshay Pai, Beomhee Park, Mathias Perslev, Ramin Rezaiifar, Oliver Rippel, Ignacio Sarasua, Wei Shen, Jaemin Son, Christian Wachinger, Liansheng Wang, Yan Wang, Yingda Xia, Daguang Xu, Zhanwei Xu, Yefeng Zheng, Amber L. Simpson, Lena Maier-Hein, and M. Jorge Cardoso. The medical segmentation decathlon. *Nature Communications*, 13(1), July 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-30695-9. URL <http://dx.doi.org/10.1038/s41467-022-30695-9>.
- Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017.
- Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005.
- Tom Boeken, Nadia Moussa, Simon Pernot, Abdellahi Abed, Carole Dean, Julien Taieb, Marc Sapoval, and Olivier Pellerin. Does bead size affect patient outcome in irinotecan-loaded beads chemoembolization plus systemic chemotherapy regimens for liver-dominant colorectal cancer? results of an observational study. *CardioVascular and Interventional Radiology*, 43(6):866–874, February 2020. ISSN 1432-086X. doi: 10.1007/s00270-020-02438-0. URL <http://dx.doi.org/10.1007/s00270-020-02438-0>.
- Tom Boeken, Claire Gallois, Richard Douard, Tchao Méatchi, Nicolas Martelli, Marc Sapoval, Julien Taieb, and Olivier Pellerin. Bridge to surgery after irinotecan-based liver chemoembolization for metastatic gastric adenocarcinoma: Letter to the editor. *Clinics and Research in Hepatology and Gastroenterology*, 45(5):101557, September 2021. ISSN 2210-7401. doi: 10.1016/j.clinre.2020.08.018. URL <http://dx.doi.org/10.1016/j.clinre.2020.08.018>.
- Adrian V. Dalca, Guha Balakrishnan, John Gutttag, and Mert R. Sabuncu. *Unsupervised Learning for Fast Probabilistic Diffeomorphic Registration*, page 729–738. Springer International Publishing, 2018. ISBN 9783030009281. doi: 10.1007/978-3-030-00928-1_82. URL http://dx.doi.org/10.1007/978-3-030-00928-1_82.
- Paul Dupuis, Ulf Grenander, and Michael I Miller. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of applied mathematics*, pages 587–600, 1998.

- Stanley Durrleman. *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. PhD thesis, 03 2010.
- Stanley Durrleman, Pierre Fillard, Xavier Pennec, Alain Trouvé, and Nicholas Ayache. Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage*, 55(3):1073–1090, April 2011. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2010.11.056. URL <http://dx.doi.org/10.1016/j.neuroimage.2010.11.056>.
- Stanley Durrleman, Stéphanie Allasonnière, and Sarang Joshi. Sparse adaptive parameterization of variability in image ensembles. *International Journal of Computer Vision*, 101(1):161–183, August 2012. ISSN 1573-1405. doi: 10.1007/s11263-012-0556-1. URL <http://dx.doi.org/10.1007/s11263-012-0556-1>.
- Stanley Durrleman, Marcel Prastawa, Nicolas Charon, Julie R. Korenberg, Sarang Joshi, Guido Gerig, and Alain Trouvé. Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage*, 101:35–49, November 2014. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2014.06.043. URL <http://dx.doi.org/10.1016/j.neuroimage.2014.06.043>.
- Théo Estienne, Marvin Lrousseau, Maria Vakalopoulou, Emilie Alvarez Andres, Enzo Battistella, Alexandre Carré, Siddhartha Chandra, Stergios Christodoulidis, Mihir Sahasrabudhe, Roger Sun, Charlotte Robert, Hugues Talbot, Nikos Paragios, and Eric Deutsch. Deep learning-based concurrent brain registration and tumor segmentation. *Frontiers in Computational Neuroscience*, 14, March 2020. ISSN 1662-5188. doi: 10.3389/fncom.2020.00017. URL <http://dx.doi.org/10.3389/fncom.2020.00017>.
- Jean Feydy. Analyse de données géométriques, au delà des convolutions, Jul 2020. URL <https://www.theses.fr/2020UPASN017>.
- Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R. Roth, and Daguang Xu. *Swin UNETR: Swin Transformers for Semantic Segmentation of Brain Tumors in MRI Images*, page 272–284. Springer International Publishing, 2022. ISBN 9783031089992. doi: 10.1007/978-3-031-08999-2_22. URL http://dx.doi.org/10.1007/978-3-031-08999-2_22.
- Alessa Hering, Lasse Hansen, Tony C. W. Mok, Albert C. S. Chung, Hanna Siebert, Stephanie Hager, Annkristin Lange, Sven Kuckertz, Stefan Heldmann, Wei Shao, Sulaiman Vesal, Mirabela Rusu, Geoffrey Sonn, Theo Estienne, Maria Vakalopoulou, Luyi Han, Yunzhi Huang, Pew-Thian Yap, Mikael Brudfors, Yael Balbastre, Samuel Joutard, Marc Modat, Gal Lifshitz, Dan Raviv, Jinxin Lv, Qiang Li, Vincent Jaouen, Dimitris Visvikis, Constance Fourcade, Mathieu Rubeaux, Wentao Pan, Zhe Xu, Bailiang Jian, Francesca De Benetti, Marek Wodzinski, Niklas Gunnarsson, Jens Sjolund, Daniel Grzech, Huaqi Qiu, Zeju Li, Alexander Thorley, Jinming Duan, Christoph Grosbrohmer, Andrew Hoopes, Ingerid Reinertsen, Yiming Xiao, Bennett Landman, Yuankai Huo, Keelin Murphy, Nikolas Lessmann, Bram van Ginneken, Adrian V. Dalca, and Mattias P. Heinrich. Learn2reg: Comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging*, 42(3):697–712, March 2023. ISSN 1558-254X. doi: 10.1109/tmi.2022.3213983. URL <http://dx.doi.org/10.1109/TMI.2022.3213983>.
- Ziyan Huang, Haoyu Wang, Zhongying Deng, Jin Ye, Yanzhou Su, Hui Sun, Junjun He, Yun Gu, Lixu Gu, Shaoting Zhang, and Yu Qiao. Stu-net: Scalable and transferable medical image segmentation models empowered by large-scale supervised pre-training, 2023.

- Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, December 2020. ISSN 1548-7105. doi: 10.1038/s41592-020-01008-z. URL <http://dx.doi.org/10.1038/s41592-020-01008-z>.
- Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics*, 30(5):509–541, 1977.
- Alice Kedra, Tom Boeken, Alessandro Di Gaeta, Charles Querub, Marc Al Ahmar, Carole Déan, Marc Sapoval, and Olivier Pellerin. Exploring a novel technique to tackle the shortage of devices for hepatic arterial infusion chemotherapy: Early results of an alternate approach for percutaneous arterial port catheter placement. *Cancers*, 15(19):4730, September 2023. ISSN 2072-6694. doi: 10.3390/cancers15194730. URL <http://dx.doi.org/10.3390/cancers15194730>.
- Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1), January 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-44824-z. URL <http://dx.doi.org/10.1038/s41467-024-44824-z>.
- Stephen Marsland and Stefan Sommer. *Riemannian geometry on shapes and diffeomorphisms*, page 135–167. Elsevier, 2020. doi: 10.1016/b978-0-12-814725-2.00011-x. URL <http://dx.doi.org/10.1016/B978-0-12-814725-2.00011-x>.
- Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- Olivier Pellerin, Tom Boeken, and Boris Guiu. Chemoembolization of hcc: Time for technical standardization, or is it too late? *CardioVascular and Interventional Radiology*, 42(9):1329–1330, June 2019. ISSN 1432-086X. doi: 10.1007/s00270-019-02268-9. URL <http://dx.doi.org/10.1007/s00270-019-02268-9>.
- Simon Pernot, Olivier Pellerin, Pascal Artru, Carole Montérymard, Denis Smith, Jean-Luc Raoul, Christelle De La Fouchardière, Laetitia Dahan, Rosine Guimbaud, David Sefrioui, et al. Intra-arterial hepatic beads loaded with irinotecan (debiri) with mfolfox6 in unresectable liver metastases from colorectal cancer: a phase 2 study. *British Journal of Cancer*, 123(4):518–524, 2020.
- Blaine Rister, Darvin Yi, Kaushik Shivakumar, Tomomi Nobashi, and Daniel L. Rubin. Ct-org, a new dataset for multiple organ segmentation in computed tomography. *Scientific Data*, 7(1), November 2020. ISSN 2052-4463. doi: 10.1038/s41597-020-00715-8. URL <http://dx.doi.org/10.1038/s41597-020-00715-8>.
- Jean-Baptiste Schiratti, Stéphanie Allasonniere, Olivier Colliot, and Stanley Durrleman. Learning spatiotemporal trajectories from manifold-valued longitudinal data. In *Advances in Neural Information Processing Systems*, pages 2404–2412, 2015.
- Alain Trouvé. An infinite dimensional group approach for physics based models in pattern recognition. *preprint*, 1995.

Jakob Wasserthal, Hanns-Christian Breit, Manfred T. Meyer, Maurice Pradella, Daniel Hinck, Alexander W. Sauter, Tobias Heye, Daniel T. Boll, Joshy Cyriac, Shan Yang, Michael Bach, and Martin Segeroth. Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence*, 5(5), September 2023. ISSN 2638-6100. doi: 10.1148/ryai.230024. URL <http://dx.doi.org/10.1148/ryai.230024>.

Laurent Younes. *Shapes and diffeomorphisms*, volume 171. Springer, 2010.

Part Three: Atlas Generation Based On Biological Data: Proof-Of-Concept Clinical Study Using Synthetic Data

August 15, 2024

Contents

1	Including biological data?	2
2	Why use ctDNA in thermal ablation	5
3	Materials and methods for assessing ctDNA	7
3.1	Patient selection, thermal ablation and follow-up	7
3.2	Next-generation sequencing of circulating tumor DNA and tumor molecular analysis	7
3.3	Outcomes	8
4	Results	9
4.1	Patient characteristics and interventions	9
4.2	Circulating tumor DNA analysis	9
5	Discussion to the prospective study	14
6	Why we need data augmentation and clustering	16
6.1	Generating Synthetic Data Using Generative Adversarial Networks	18
6.2	Modelling Atlases Using the MCMC-SAEM Algorithm	18
6.3	Discussion	22
7	REFERENCES	24




Figure 1: PhD map.

This section refers to the *orange projects*.

This section was adapted from the following publications:

Clinical value of sequential circulating tumor DNA analysis using next generation sequencing and epigenetic modifications for guiding thermal ablation for colorectal cancer metastases: a prospective study. Submitted Lancet Gastroenterology & Hepatology

Boeken T, Blons H, Dean C, Sapoval M, Pellerin O. The potential of ctDNA in locoregional therapies for colorectal cancer. Lancet Gastroenterol Hepatol.

Atlas Generation from Synthetic Data for Predicting Trajectories of Circulating Tumor DNA in Thermal Ablation Referrals: Proof-of-concept study. To be submitted

Special thanks for this section: Professor H el ene Blons for the ongoing support in oncological genetics, Cl ement Mantoux and Juliette Chevallier for their help in modelling circulating DNA dynamics.

1 Including biological data?

A novel *biological* biomarker emerged as a breakthrough for the management of colorectal cancer: **circulating tumor DNA (ctDNA)**.

Circulating tumor DNA (ctDNA) consists of DNA from tumor cells circulating in the blood. Detected in minute quantities through sequencing techniques, this biomarker allows for more sensitive and early estimation of tumor burden compared to conventional imaging methods. Its prognostic relevance has been demonstrated not only in non-metastatic colorectal disease but also in selecting patients for surgery in oligometastatic disease Boysen et al. (2020)

Our **prospective ancillary project** to this thesis aimed at clarifying the prognostic and therapeutic value of repeated analysis of circulating tumor DNA before and after treatments by interventional radiologists, such as ablation. **We decided to collect and analyze samples from our patients, and we believe to be the first to use ctDNA in such setting.**

Interventional radiology provides a non-surgical alternative through percutaneous treatment. In this project, only thermal ablation (using microwave or radiofrequency for heat-based destruction) is considered. Data on hepatic metastases thermal ablation, driven by major randomized trials

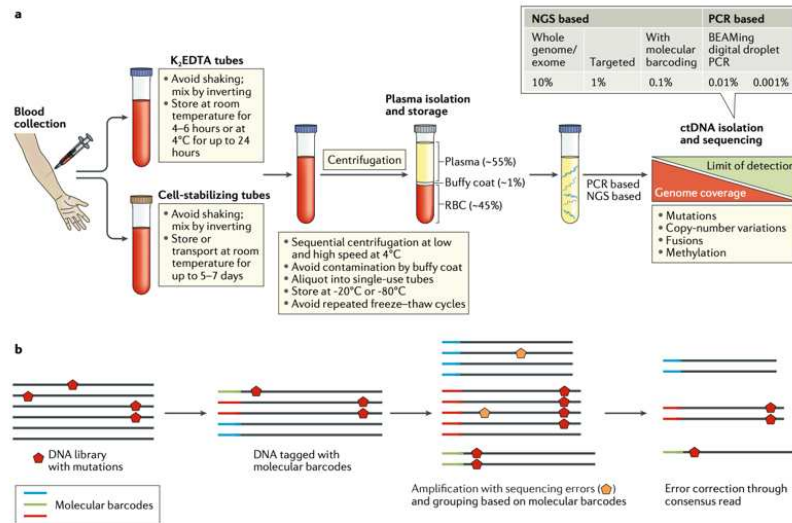


Fig. 2 | ctDNA isolation and analysis. a | Circulating cell-free DNA, of which circulating tumour DNA (ctDNA) is a part of, is isolated from plasma samples after serial centrifugation of blood collected in either K₂EDTA or cell-stabilizing tubes. ctDNA is subsequently isolated from cell-free DNA using library preparations and analysed for the presence of various genetic aberrations, including mutations, copy-number variations, fusions and/or other alterations such as changes in DNA methylation. **b** | Molecular barcoding: prior to PCR and sequencing, each DNA fragment can be labelled with unique DNA barcodes; subsequently, reads that share the same barcode (typically in thousands) can be grouped together because they all originate from the same ctDNA fragment. This approach enables sequencing errors (orange pentagon, seen in the minority) to be distinguished from true mutations (red pentagon, seen in the majority). Molecular barcoding also helps correct potential biases in amplification and thus assists in the precise quantification of mutations or amplification frequencies. NGS, next-generation sequencing; RBC, red blood cells.

Figure 2: Circulating DNA collection, from Dasari et al.

such as the COLLISION trial Puijk et al. (2018), have positioned this technique in the therapeutic arsenal. They showed increased overall survival in patients treated with thermal ablation compared to systemic treatment alone. This treatment delayed progression and cured a fraction of patients, although half of the treated patients ultimately succumbed to their disease within 5 years of treatment. In comparison, after hepatic surgery, 70% of treated patients experience recurrence or progression within two years. The optimal place of curative-intent thermal ablation in the patient's oncological journey, however, remains to be defined, especially in relation to overall disease control through systemic treatments.

The oncology and biology teams at HEGP have been experts in the field of ctDNA for many years, particularly through the Carpem project. This technique exhaustively searches for molecular anomalies related to the studied genes. The panels target disease-specific genes; their reasonable cost and rapid results enable routine clinical use at HEGP. Indeed, ctDNA is routinely measured at HEGP in cases with ambiguous disease progression or stability.

Such a biomarker would be particularly suited for the clinical situation of patients undergoing hepatic thermal ablation. New biomarkers—circulating tumor DNA, circulating tumor cells, circulating extracellular vesicles, etc.—have introduced the concept of minimal residual disease (MRD), and more broadly, new definitions of oligometastatic disease. Beyond semantics, these new clinical entities have therapeutic implications. Based on the stage (single metastasis to multi-metastatic disease), treatment options evolve along with advancements in locally ablative techniques in the liver (thermal ablation, electroporation, radiotherapy, intra-arterial therapies, and minimally invasive




Figure 3: Stages definitions, from Dasari et al.

surgery).

Circulating tumor DNA (ctDNA) has been studied in several contexts throughout the oncological history of colorectal cancer. In non-metastatic colorectal disease, publications agree on the prognostic value of the biomarker Loft et al. (2023): patients with positive ctDNA before or after surgery exhibited worse prognoses than patients without circulating tumor DNA. The presence of ctDNA after surgery reflects residual disease and places the patient in a high-risk group for early recurrence. Similarly, in metastatic disease, persistent circulating tumor DNA reflects inadequately treated disease. It is detected in over 90% of cases.

Naturally, *surgical* teams have examined ctDNA in the perioperative setting or in potentially resectable disease. The ancillary study of the Unicancer Prodiges-14 trial Bidard et al. (2019) focused on 153 patients with potentially resectable disease. Among patients who underwent R0/R1 resection, those with ctDNA+ before surgery had significantly shorter overall survival. The study concluded on the utility of ctDNA for patient selection before surgery in oligometastatic disease. In 2019, the team at Memorial Sloan Kettering Cancer Center conducted a study on 60 patients treated with curative-intent hepatic surgery Narayan et al. (2019). Two-year progression-free survival was 30%, and progression-free survival was associated with perioperative ctDNA detection. The Melbourne team similarly utilized this principle in a 2021 prospective study involving 54 patients and 380 pre- and postoperative samples from hepatic surgery patients Tie et al. (2021). All patients had upfront resectable disease and were treatment-naïve at inclusion. ctDNA was detected in 85% of patients at inclusion and 24% after surgery. The study detailed the importance of ctDNA clearance (i.e., disappearance of detection under neo- or adjuvant systemic treatment). Postoperative ctDNA+ patients were at higher risk of recurrence: (HR 6.3; 95% CI 2.58 to 15.2; $P < 0.001$). Overall survival was also strongly associated with postoperative detectability: HR 4.2; 95% CI 1.5 to 11.8; $P < 0.001$. This study also highlighted the value of serial ctDNA measurements after surgery, particularly during adjuvant treatment. All patients with postoperative ctDNA+ after adjuvant chemotherapy had experienced recurrence during follow-up. Patients with postoperative ctDNA+ that subsequently became negative during adjuvant chemotherapy had a




Figure 1: The origin of ctDNA, cancer-related changes in ctDNA, and factors affecting ctDNA concentrations and detection
 ctDNA=cell free DNA. CTC=circulating tumour cells. ctDNA=circulating tumour DNA. *The greater-than sign represents differences in ctDNA concentrations by various factors.

Figure 4: Clinical applications of ctDNA in colorectal cancer, from Loft et al.

better prognosis, suggesting that this biomarker could also reflect the effectiveness of postoperative systemic treatment. This aligns with results from non-metastatic cancers and the role of adjuvant therapies, tailored based on individual patient risk.

The results from these surgical series cannot be extrapolated to thermal ablation due to differing biological and molecular mechanisms, even though the ultimate goal remains tumor clearance (i.e., complete reduction of tumor burden and negative ctDNA detection). During thermal ablation, the tumor remains in place, and the thermal energy induces necrosis, ischemia, apoptosis, and/or immunogenic cell death, generating intrinsic immunogenicity (priming concept).

Published data on thermal ablation and ctDNA are very limited. A preliminary study including patients treated with thermal ablation and radiotherapy showed that all five patients with postoperative ctDNA+ had a worse prognosis.

As of now, no publication has addressed different ctDNA evolutionary profiles following thermal ablation. Thus, we decided to add this feature to our work on patient selection.

2 Why use ctDNA in thermal ablation

The field of precision medicine for patients with colorectal cancer metastatic disease is rapidly expanding. Molecular profiling is now recommended by the latest guidelines to adequately select systemic therapies Cervantes et al. (2023). Profiling for *locoregional therapies*, however, has not yet

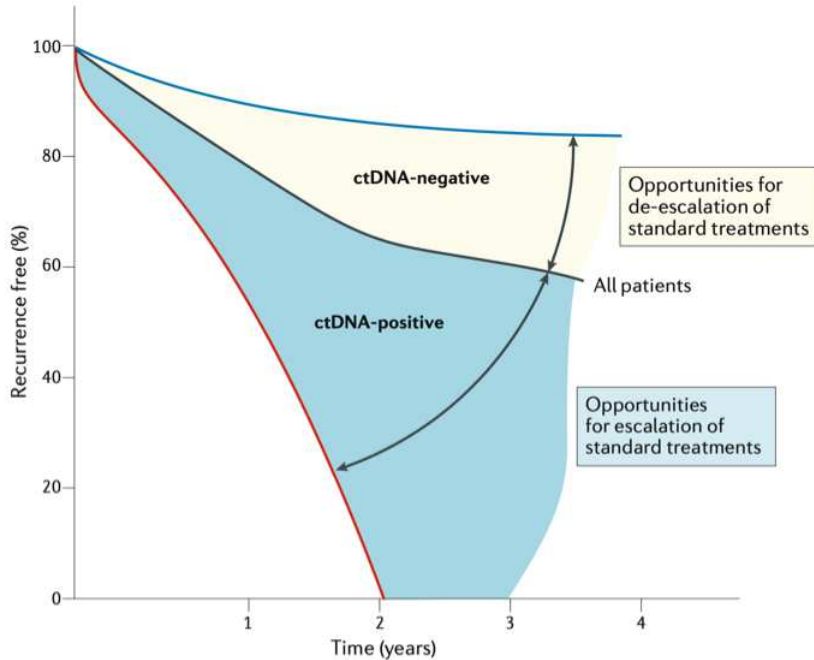


Figure 5: Opportunities as a biomarker, from Dasari et al.

been evaluated in practice. The clinical value of performing sequential assessments of circulating tumor DNA (ctDNA) for patients who are referred to thermal ablation is not known.

Alongside the progress made in systemic treatments, interventional radiologists have customized locoregional solutions for patients diagnosed with oligometastatic disease. Thermal ablation, when technically feasible, has consistently shown to be a valid treatment option for patients with oligometastatic disease, particularly in lung and liver target lesions. Despite these advances, the reported 5-year overall survival of thermal ablation for colorectal metastases remains below 50%, even when using the latest assisting technologies De. These results highlight the need for further research focused on enhancing patient selection and monitoring responses to therapy.

Our team developed an accurate method to detect tumor mutations in plasma from Next-Generation Sequencing (NGS) data based on the quantification of the error rate of each base position using the panel (AmpliSeq™ Colon Lung V2 Panel, BPER) Pécuchet et al. (2016). This algorithm was first validated in lung cancer and was successfully used in a prospective trial validating plasma testing for routine analysis in colorectal cancer patients. For tumors that do not have mutations in genes tested by the panel, we analyze plasma samples by using the detection of two methylated markers (WIF1 and NPY) by digital droplet PCR accordingly to a method developed and validated for colorectal cancer. Such liquid biopsies have the potential to serve as a surrogate for tumoral burden and minimal residual disease (MRD), aiding in the evaluation of recurrence risk following thermal ablation.

These methods are now available for use in routine care Moati et al. (2018), as demonstrated by specialized centers, and can be used to guide therapies in colorectal cancer. As an example, a recent significant randomized controlled trial demonstrated the successful ability of ctDNA to

differentiate patients who would benefit from adjuvant therapies following the removal of stage II colorectal cancer Tie et al. (2022). The ctDNA biomarker-driven prospective approach was non-inferior to the standard approach, and could set a new standard of care for patients with stage II colorectal cancer.

Unfortunately, the dynamics of ctDNA before and after thermal ablation are not known. The serial analysis of ctDNA in thermal ablation could help understand the kinetics of tumor ctDNA release and the impact of therapeutic thermal damage on ctDNA; more importantly, it could help improve the overall patient management throughout the course of disease.

The hypothesized result is that ctDNA monitoring predicts early post thermal ablation relapse.

This prospective study was therefore designed with the aim of determining the prognostic value of ctDNA before thermal ablation in terms of recurrence-free survival using next-generation sequencing (NGS) combined with the assessment of epigenetic modifications.

3 Materials and methods for assessing ctDNA

3.1 Patient selection, thermal ablation and follow-up

This is a single-center single arm observational prospective study that was approved by IRB #00011928 (CERAPHP Centre 2023-01-05). All consecutive patients treated between November 2021 and June 2022 with colorectal cancer liver, lung or bone metastases referred by the multidisciplinary tumor board to curative-intent thermal ablation were screened.

Inclusion criteria were colorectal cancer metastases accessible to curative intent thermal ablation as assessed by the tumor board, measurable disease burden (RECIST 1.1 base method), age above 18 years, ECOG 0 or 1, stable disease under systemic treatment, feasible follow-up, resected primary colorectal tumor. Exclusion criteria were combined treatment with surgery or radiotherapy, unresected primary tumor, impaired liver, lung or kidney function, refusal to participate in the study and concomitant non-colorectal cancer.

The interventions were planned under general anesthesia during a two-day hospital stay and bevacizumab was interrupted for a minimum of three weeks before any interventional procedure. Tumor size, localization, molecular profile, primitive tumor localization, types and number of prior treatment lines, ablation margins and post-interventional treatments were recorded. Thermal ablation was microwave ablation for liver lesions or cryoablation or radiofrequency for lung or bone lesions. Adverse events were scored according to CTCAE criteria. Radiological data, including CT and MRI images before thermal ablation and at follow-up, were reviewed by an independent observer. Routine clinical, biological and imaging were performed. Follow-up included a radiological reevaluation at one month, two months, and every three months afterwards in collaboration with the oncological team.

3.2 Next-generation sequencing of circulating tumor DNA and tumor molecular analysis

Tumor molecular profiles were assessed as recommended for metastatic colorectal cancer in care settings. Briefly, tumors are analyzed using a 29-genes targeted NGS panel including KRAS, NRAS, BRAF, TP53, PIK3CA, CTNNB1, ERBB2 for the most frequently mutated in colorectal cancer.




Figure 6: Study workflow: NGS refers to Next Generation Sequencing of circulating DNA and WIF and NPY refer to the assessment of a hypermethylation of WNT inhibitory factor 1 (WIF1) and neuropeptide T (NPY) using droplet PCR.

Tumor molecular profiles were known for 80% of the patients referred to thermal ablation, and were unknown for 20% of the patients because they were referred from another center.

Blood collection was performed at the induction of the anesthesia (Hour 0 or H0), the emergence (H2), the morning following the intervention (H24) and the first follow-up consultation (usually on Day 30 or D30). Additional samples were analyzed according to the oncologist's decision, as represented in the **Figure below**. Blood collection was performed in a 10 mL blood Collection EDTA tube. Upon receipt, the tubes were centrifuged 10 minutes at 1600 g, the plasma recovered and a second centrifugation was carried out for 10 minutes at 6000 g. The plasma was then transferred in LoBind tubes (Eppendorf) and stored at 80°C until ctDNA analysis. Time from collection to storage was less than 6 hours as recommended in routine care settings in our institution.

ctDNA was extracted using the Maxwell[®] RSC ccfDNA Kit (Promega), quantified by Qubit dsDNA HS Assay Kit. cfDNA was analyzed using the AmpliSeq[™] Colon Lung Cancer Panel (ThermoFisher) and libraries were generated automatically on the ionchef (AmpliSeq[™] Kit for Chef DL8) from 15ul of cfDNA. NGS was performed using the Ion Proton with minimal depth of 3000 X. Sequencing data were analyzed using the BPER method as previously described. For patients with no available data on tumor tissues or wild type tumors, negative plasma samples were assessed by WIF and NPY methylation-ddPCR (met-ddPCR) to validate the absence of ctDNA as previously described.

All presented mutations are known pathological tumor mutations. If multiple mutations were identified, the most prevalent one was selected for subsequent statistical analysis.

3.3 Outcomes

The primary endpoint was recurrence-free survival, defined as the time from the intervention to disease recurrence, according to RECIST version 1.1 criteria, or death (or the date of the last follow-up if the patient was alive with no recurrence). Secondary endpoints were recurrence at any

site (defined as the time from intervention until the first radiological recurrence at any site), and local recurrence (defined as the time from intervention until the first radiological recurrence on a targeted site).

All analyses are reported according to the STROBE statement. No sample size and power calculation were done. Sample size was chosen based on the estimated enrolment rate and an acceptable study period. Continuous variables were presented as means (\pm SD) for normally distributed data and otherwise as median and interquartile range (IQR). Categorical data are numbers and percentages. No replacement of missing data was performed.

Recurrence-free survival were analyzed with a two-sided log-rank test, with the hazard ratio (HR) and two-sided 95% CIs based on a Cox proportional-hazards model and the associated Kaplan-Meier survival estimates. Patients alive and recurrence-free were censored on the date of the last evaluable tumor assessment. The proportional-hazards assumption was assessed by testing the interaction between covariates and time. Predominant clonal ratio pre-op (%) was dichotomized according to its median value ($=0\%$ vs $>0\%$).

4 Results

4.1 Patient characteristics and interventions

A total of 19 consecutive patients were screened between November 2021 and June 2022. Four patients were not included in this study because of non-intent to cure thermal ablation (these included three patients with primary tumor in place and one patient awaiting portal embolization). Fifteen patients (mean age: 65.3 ± 13.4 years) were included and analyzed in this study.

Baseline patient characteristics are summarized in **Table 1**. All patients were previously treated by systemic therapies, with a median number of lines of 1 [range 1-2]. Thermal ablation was performed in a single session and targeted a median number of lesions of 1 [range 1-3]. Target lesions were respectively located in the liver (for $N = 12/15$ patients, using microwave), the lung (for $N = 2/15$ patients, using cryotherapy for one patient and radiofrequency for one patient), and within the iliac bone (for $N = 1/15$ patient, using cryotherapy). The mean diameter of the largest target was $19.2 \text{ mm} \pm 14 \text{ mm}$. Tumor characteristics are available in **Table 2**.

All thermal ablations were technically successful as evaluated by the interventional radiologist. No early or late complication from the fifteen interventions were recorded and all patients were discharged on the following day.

The median follow-up was 316 days. By the end of the study, 6 patients ($N = 6/15$, 40%) remained disease free and 9 patients ($N = 9/15$, 60%) had recurred. The recurrence was both local (on the treated site) and distant for 2 patients ($N = 2/15$, 13%) on day 29 for the first patient and day 49 for the second patient. For 7 patients (46%) recurrence was exclusively distant ($N = 7/15$, 46%). No mortality was reported during follow-up. No patient had isolated local recurrence on the targeted lesion. The median recurrence-free survival was 250 days.

4.2 Circulating tumor DNA analysis

For patients treated in our institution, tumor samples were available and analyzed by NGS. Molecular profiles were concordant between plasma and tumor for patients with positive ctDNA. For 7/10 patients referred to our center for thermal ablation only, information on tumor molecular profiles were available in medical records otherwise reported unknown. Five tumors were reported

Sex	F/M	6/9 (40.0%/60%)
Age at intervention (years)		65.3 ± 13.4
Primary tumor localization	Hinge	2 (13.3%)
	Left colon	1 (6.7%)
	Rectum	4 (26.7%)
	Right colon	4 (26.7%)
	Sigmoid	4 (26.7%)
Ablation site	Bone	1 (6.7%)
	Liver	12 (80.0%)
	Lung	2 (13.3%)
ECOG status	0	8 (53.3%)
	1	7 (46.7%)
Number of lines prior to treatment	1	8 (53.3%)
	2	7 (46.7%)
First treatment line	FOLFIRINOX BEVACIZUMAB	1 (6.7%)
	FOLFOX	10 (66.7%)
	FOLFOX BEVACIZUMAB	1 (6.7%)
	FOLFOX CETUXIMAB	1 (6.7%)
	NA	2 (13.3%)
Treatment at intervention date	FOLFIRI	5 (33.3%)
	FOLFIRI BEVACIZUMAB	2 (13.3%)
	FOLFIRI CETUXIMAB	1 (6.7%)
	FOLFOX	5 (33.3%)
	NA	2 (13.3%)
Number of ablation sites	1	10 (66.7%)
	2	4 (26.7%)
	3	1 (6.7%)
Intent to cure		15 (100.0%)
Largest target diameter (mm)		19.2 ± 14.0
Ablation technology	Cryotherapy	2 (13.3%)
	Microwave	12 (80.0%)
	Radiofrequency	1 (6.7%)
CTDNA samples per patient		4.0 [3.0;4.0]

Figure 7: **Table 1.** Baseline characteristics. Continuous variables were presented as means (\pm SD) for normally distributed data and otherwise as median and interquartile range (IQR). Categorical data are numbers and percentages. No replacement of missing data was performed.

Tumor mutations determined by NGS in our Institution*(N=5) Information in medical records tissues not available for centralized testing (N=1)***		Plasma mutation (NGS)	Variant Allele Fraction	Predominant clonal ratio: plasma tumor DNA (ng/mL)	Methylation (WIF/NPY)	Circulating tumor DNA status (Negative/Positive)
Patient 1	KRAS Mutated**					
- H0		0	0%	0	NP	N
- H2		0	0%	0	NP	N
- H24		KRAS p.Gly12Ala c.35G>C	1.8%	0.18	NP	P
Patient 2	TP53 (NM_000546.5) p.Arg248Trp c.742C>T; TP53 (NM_000546.5) p.Arg158His c.473G>A*	TP53 p.Arg248Trp c.742C>T TP53 p.Arg248Trp c.742C>T / TP53 p.Arg158His c.473G>A TP53 p.Arg248Trp c.742C>T / TP53 p.Arg158His c.473G>A	1.0% 2.3% / 2.1% 1.4% / 1.2%	0.15 0.81 3.16	NP NP NP	P P P
Patient 3	KRAS & BRAF V600E WT**	PIK3CA p.His1047Arg c.3140A>G/ KRAS p.Gly12Val c.35G>T 0 PIK3CA p.His1047Arg c.3140A>G/ KRAS p.Gly12Val c.35G>T	0.2% / 0.2% 0 0.5% / 0.15%	0.15 0 0.35	NP NP NP	P N P
Patient 4	KRAS Mutated**					
- H0		0	0%	0	N	N
- H2		0	0%	0	N	N
- H24		0	0%	0	N	N
Patient 5	KRAS Mutated**					
- H0		0	0%	0	N	N
- H2		0	0%	0	N	N
- H24		0	0%	0	2.7 / 1.14	P
Patient 6	KRAS WT**					
- H0		0	0%	0	NP	N
- H2		0	0%	0	NP	N
- H24		BRAF p.Asp594Gly c.1781A>G 5 / TP53 p.Cys242Tyr c.725G>A	1% / 1%	0.09	NP	P
Patient 7	KRAS (NM_033360.3) p.Gly12Val c.35G>T/PIK3CA (NM_006218.2) p.His1047Arg c.3140A>G TP53 (NM_000546.5) p.Arg175His c.524G>A*	KRAS p.Gly12Val c.35G>T / PIK3CA p.His1047Arg c.3140A>G / TP53 p.Arg175His c.524G>A KRAS p.Gly12Val c.35G>T / PIK3CA p.His1047Arg c.3140A>G / TP53 p.Arg175His c.524G>A NA	3.6% / 1.7% / 1.5% 1.6% / 1% / 1.6% NA	0.31 0.98 NA	NP NP NA	P P NA
Patient 8	Unknown tumor status (ctDNA tested by WIF NPY methylation)**					
- H0		0	0%	0	N	N
- H2		0	0%	0	N	N
- H24		0	0%	0	N	N
Patient 9	T1: PIK3CA (NM_006218.2) p.Glu542Lys c.1624G>A TP53 (NM_000546.5) p.Arg306Ter c.916C>T * T2: CTNNB1 (NM_001904.3) p.Thr41Ala c.121A>G TP53 (NM_000546.5) p.Arg175His c.524G>A *					
- H0		0	0%	0	NP	N
- H2		0	0%	0	NP	N
- H24		0	0%	0	NP	N
Patient 10	KRAS (NM_033360.3) p.Gly12Asp c.35G>A SMAD4 (NM_005359.5) p.Asp355His c.1063G>C*					
- H0		0	0%	0	NP	N
- H2		0	0%	0	NP	N
- H24		0	0%	0	NP	N
Patient 11	Unknown tumor status					

Mutations involved the following genes: KRAS (NM_033360.3), PIK3CA (NM_006218.2), BRAF (NM_004333.4) and TP53 (NM_000546.5)

Methylation assessment is indicated: WIF and NPY refer to the assessment of a hypermethylation of WNT inhibitory factor 1 (WIF1) and neuropeptide T (NPY) using droplet PCR.

Patient 10 had 2 tumors designed T1 and T2 with different molecular profiles.

NA: missing data or unavailable

NP: not performed

N: negative

P: positive

NGS: next-generation sequencing

Mutations involved the following genes: KRAS (NM_033360.3), PIK3CA (NM_006218.2), BRAF (NM_004333.4) and TP53 (NM_000546.5)

Methylation assessment is indicated: WIF and NPY refer to the assessment of a hypermethylation of WNT inhibitory factor 1 (WIF1) and neuropeptide T (NPY) using droplet PCR.

Patient 10 had 2 tumors designed T1 and T2 with different molecular profiles.

NA: missing data or unavailable

NP: not performed

N: negative

P: positive

NGS: next-generation sequencing

-	H2		0	0%	0	N	N
-	H24		0	0%	0	N	N
Patient 12		KRAS mutated**					
-	H0		0	0%	0	N	N
-	H2		BRAF p.(Asn581Thr) c.1742A>C	0.8%	0.12	0.37 / 0	P
-	H24		0	0	0	N	N
Patient 13		Unknown tumor status but mutations in KRAS and TP53 are expected**					
-	H0		KRAS p.Gly125Ser c.34G>A TP53	14% /	2.38	NP	P
-	H2		p.Pro152Thr c.454C>A	19%	NA	NP	NA
-	H24		NA	NA	3.1	NP	P
			KRAS p.Gly13Asp c.38G>A/ TP53	14.6% /			
			p.Pro152Thr c.454C>A	19.3%			
Patient 14		KRAS mutated**					
-	H0		0	0%	0	NP	N
-	H2		PIK3CA p.Gly542Lys c.1624G>A	1.2%	0.001	NP	P
-	H24		0	0	0	NP	N
Patient 15		KRAS (NM_033360.3) p.Gly12Ala c.35G>C TP53 (NM_000546.5) p.Arg282Trp c.844C>T*					
-	H0		0	0%	0	NP	N
-	H2		KRAS p.Gly12Ala c.35GG>T/ TP53	2.3% /	0.9	NP	P
-	H24		p.Arg282Trp c.844C>T	5.0%	0	NP	N
			0	0%			

Figure 9: Table 2 (continued)

KRAS mutated. For one patient, the *KRAS* mutation was detected in plasma, one patient had a *BRAF* p.N581T mutation in plasma, one patient a *PIK3CA* mutation in plasma, one patient had no ctDNA, and one was positive based on methylation only. One tumor was reported *KRAS* WT, this patient had a *TP53* & a *BRAF* p. D594G mutation, one tumor was reported *KRAS* & *BRAF* WT and was *KRAS* and *PIK3CA* positive in plasma, as described in **Table 2**.

A median number of 4 [range: 3-9] successful DNA extractions were performed per patient, resulting in the analysis of a total of 60 samples. Pathological mutations were detected for 9 patients during the course of the study ($N = 9/15$, 60%) in 34 samples ($N = 34/60$, 57%). Mutations were identified in the following gene sequences: *KRAS*, *PIK3CA*, *TP53*, *BRAF* and *SMAD4*. Circulating tumor DNA characteristics are provided in **Table 2**. The assessment of a hypermethylation of WNT inhibitory factor 1 (WIF1) and neuropeptide T (NPY) using droplet PCR was positive for two with NGS-negative ctDNA, respectively with WIF/ NPY ratios of 2.71%/1.14% and 0.35%/0%. An example is provided in the **Figure below**.

ctDNA was positive for 33% of the samples collected during the first 24 hours. The prognosis value of baseline ctDNA was analyzed as follows. Patients were defined as ctDNA positive if ctDNA was detected prior to the intervention. The ratio of ctDNA positive / negative patients at baseline was 27% / 73%. The mean total plasma DNA concentration prior to the intervention was 15.53 ± 19.00 ng/mL and the mean mutation ratio of the predominant mutation when detected was 4.70 ± 6.37 . Up to three DNA mutations were simultaneously detected in the samples. The median recurrence-free survival was respectively for the ctDNA+ population 46 days (IC 95%: 29-not reached) and for the ctDNA- population not reached (CI 95%: 51-not reached) ($p = 0.038$). The hazard ratio for progression according to the presence of baseline circulating tumor DNA was estimated at 0.14 (CI 95%: 0.03-0.65, $p = 0.019$). The survival curves are provided in the **Figure below**.

The impact of thermal ablation on the initial dynamics of ctDNA was described as follows. The dynamics of ctDNA during the first 24 hours are provided in the **Figure below**. The Spearman correlation coefficient between pre and post ablation ctDNA quantifications was estimated at 0.42

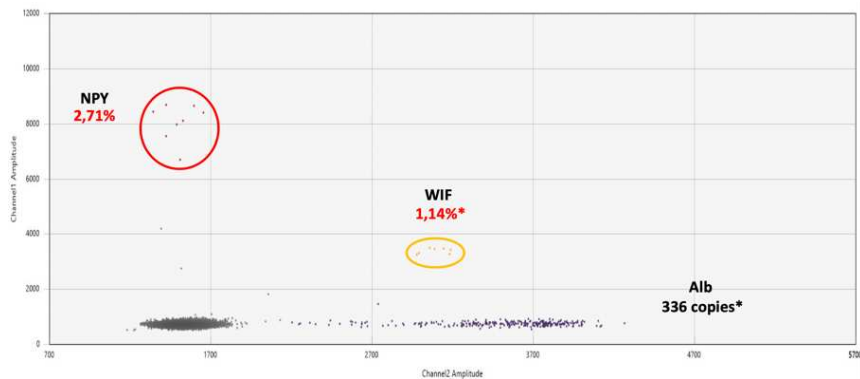


Figure 10: Example of a patient with NGS ctDNA negative samples and positive hypermethylation of WNT inhibitory factor 1 (WIF1) and neuropeptide T (NPY) using droplet PCR.




Figure 11: - Recurrence-free survival according to the baseline circulating tumor DNA. Progression-free survival were analyzed with a two-sided log-rank test, with the hazard ratio (HR) and two-sided 95% CIs based on a Cox proportional-hazards model and the associated Kaplan-Meier survival estimates. Patients alive and recurrence-free were censored on the date of the last evaluable tumor assessment.




Figure 12: Evolution of the predominant clonal ratio (%) from circulating tumor DNA as assessed before thermal ablation (H0), immediately after (H2) on the following day (H24), and up to the last blood sample available.

($p=0.13$), with 80% of patients presenting a positive change in clonal ratios between H0 and H2. ctDNA at H2 was not predictive of recurrence. At H2, the median recurrence-free survival was respectively for the ctDNA+ population 162 days (IC 95%: 49-not reached) and for the ctDNA- population 278 (CI 95%: 30-not reached) ($p = 0.92$). ctDNA at H24 was not predictive of recurrence. At H24, the median recurrence-free survival was respectively for the ctDNA+ population 103 days (IC 95%: 29-not reached) and for the ctDNA- population not reached (CI 95%: 51-not reached) ($p = 0.06$). All patients with no recurrence during follow-up were ctDNA negative at H24. At the one-month follow-up, three patients remained ctDNA positive, and all three patients recurred during the study follow-up.

5 Discussion to the prospective study

Our study demonstrated the feasibility of assessing circulating tumor DNA and epigenetic modifications in routine care for patients undergoing thermal ablation of colorectal cancer metastases. Our team developed a sensitive technique utilizing Next-Generation Sequencing of ctDNA, which successfully detected ctDNA in more than half of the patients throughout the study. For patients with no or unknown tumor mutations status we validated that NGS negative results were not a false negative using colorectal cancer specific epigenetic modifications. Two negative samples were positive for WIF and/or NYP markers showing that agnostic biomarkers are important to validate negative NGS results. Methylation markers may rescue false negative NGS tests in some patients.

Our findings are particularly noteworthy as patients referred for thermal ablation typically

exhibit minimal tumor burden and limited disease. Our study population presented with 1-3 target lesions measuring an average of 2cm. In this oligometastatic context, we used a very low technical threshold (<1%) for ctDNA detection, and the proposed method yielded a positive signal for ctDNA as a predictive biomarker for early relapse after thermal ablation.

The median recurrence-free survival was 250 days, and no patient experienced exclusive local recurrence. These results align with published data on thermal ablation as a therapy for locoregional disease control. However, a majority of patients in our study experienced recurrence, emphasizing the need for a deeper understanding of metastatic disease and the overall role of locoregional treatments in the treatment algorithm. Notably, patients with detectable ctDNA before the intervention experienced rapid recurrence compared with ctDNA negative patients. This result while based on a small population, suggests the strong potential of ctDNA as a clinically relevant biomarker. Conducting larger studies with greater statistical power may further elucidate whether patients with positive ctDNA would benefit from locoregional therapies. Moreover, achieving total extinction of ctDNA prior to thermal ablation could become an intermediary objective within systemic treatment.

While ctDNA detection at H2 did not predict early relapse, detection at H24 was found to be related to recurrence suggesting that cell death induced by thermal ablation impacts ctDNA kinetics. All patients with no recurrence during follow-up were ctDNA negative at H24, and the presence of ctDNA at the one-month follow-up was a statistically significant hazard ratio for recurrence. This suggests that achieving ctDNA extinction could also serve as an important treatment objective, in addition to achieving a complete response based on radiological assessment.

The immediate dynamics of ctDNA before and after thermal ablation bring new data to this expanding field. The initial description of cell-free DNA in human plasma was published in 1948 Mandel and Metais (1948). Since then, in vivo and in vitro studies have focused on better explaining the origins of circulating DNA in patients. Theories relate to the release of ctDNA via necrosis and apoptosis but also through oncosis (ischemic cell death), pyroptosis, phagocytosis, active secretion, neutrophil extracellular trap release (NETosis), and excision repair A et al. (2019). Conversely, while variations in ctDNA are typically attributed to the tumor burden, they might also reflect tumor metabolism. These biological features might explain why some studies show an increase of cell-free ctDNA concentrations during the initial phase of radiation or systemic therapy Cheng et al. (2009). The initial 24-hour period seems particularly prone to ctDNA dynamic changes that could help predict long-term outcomes.

We also report genetic heterogeneity between plasma and tumor tissue. Technical issues, especially for patients screened in another institution, molecular switch for patients initially KRAS wild type with KRAS positive ctDNA, second cancer or emergence of a tumor subclone are possible explanation for these discrepancies.

There is no comparable data to date focusing on this specific situation. We provided valuable additional information regarding patients with oligometastatic limited disease. Given the complexity of the cell death induced by thermal ablation, results from the analysis of ctDNA dynamics before and after surgery are not directly transposable. Thermal ablation has been shown to induce coagulation necrosis and apoptosis through both direct and indirect mechanisms, and to activate an antitumor immune response, all of which are distinct from the effects of surgical resection.

Our clinical study has several limitations. One limitation to our study was the use of recurrence-free survival as a surrogate for overall survival, a method often employed in studies regarding locoregional treatments. The underlying hypothesis is that our local treatment will delay the overall recurrence and hence improve the overall survival. A recent meta-analysis assessed the correlation between recurrence-free survival and overall survival after resection of colorectal liver metastases




Figure 13: *Three trajectories of plasma DNA concentrations from the real dataset.*

for more than 3000 patients and found a minimal correlation Ecker et al. (2022). Future studies on the role of ctDNA as a predictive factor for poor response should probably be designed with overall survival as a primary outcome. Another limitation is the study size. As a proof-of-concept feasibility study, and given the costs of the techniques used, we could not provide sufficient power for yielding statistically significant differences in all measures. Conclusions are exploratory and should only advocate further studies on the subject.

In conclusion, circulating tumor DNA may discern patients who are likely to benefit from thermal ablation from those who may not. Its clinical utility in patient selection holds great promise and warrants comprehensive investigation through extensive large-scale prospective studies.

6 Why we need data augmentation and clustering

The initial dataset is **relatively small** given the complexity of the study protocol and its implementation. We were quite thrilled by our ability to use ctDNA in routine thermal ablation, because such a biomarker was never assessed before in the peri operative setting.

Our proof-of-concept study provided valuable data regarding the preoperative prognostic value of CTDNA but *the trajectories* (in terms of disease progression throughout time) were not studied.

Can we generate an atlas for predicting clinical outcome based on trajectories of circulating tumor DNA in thermal ablation referrals?

Following our published statement in the Lancet, we propose to generate synthetic data from this proof-of-concept clinical evaluation that analyzing the genetics and epigenetics of the tumor




Figure 14: *Two clusters: dotted lines for progressive patients.*

within the plasma for patients with colorectal cancer metastases referred to thermal ablation could enhance patient selection.

This was the first clinical study conducted on thermal ablation and systematic ctDNA use, with both next generation sequencing (using a proprietary gene panel) and specific hypermethylation parameters (using ddPCR). However, because of the technicity of this analysis, and its current financial cost, we lack large datasets that could provide sufficient significance for the results.

The objective of the current section is to compute two population trajectories from circulating tumor DNA ratios, and to produce an atlas for good-outcome patients following thermal ablation and an atlas for bad-outcome patients in terms of cancer recurrence. Such atlases could then be confronted with the hypothetical kinetics of circulating tumor DNA in the peri-operative setting.

Gene mutations in cancer patient plasma throughout time can be treated as longitudinal data that assume a correlation between observations, incorporating latent variables to capture inherent variability. When using mixed-effects models, we combine fixed effects (consistent parameters across individuals, or population trends) with random effects (varying parameters represented by latent variables, or individual deviations).

The trajectory of each individual can be modeled as a piecewise affine trajectory. The population trajectory is not obvious, but the hypothesis stemming from biological data was that thermal ablation would initially release ctDNA between Hour 0 and Hour 2 immediately following the thermal ablation, then decrease gradually to reach a low point before either perpetuate at negative values (for good outcome patients), or augment in later stages (for bad outcome patients). Moreover, time breaks are unknown: we arbitrarily chose to sample patients at fixed time points (H0, H2, H24 and D30) that do not reflect the natural time breaks in terms of disease response, disease recurrence or ctDNA clearance. The present use of mixed-effects models adapted to longitudinal datasets has the potential of revealing new time breaks involved in the biological mechanisms involved in thermal ablation.

6.1 Generating Synthetic Data Using Generative Adversarial Networks

For demonstration purposes, we propose to model each individual as a piecewise affine trajectory. We can use data augmentation, generative models or resampling. One option is to use a Generative Adversarial Network (GAN) ¹⁶ with two neural networks, the Generator and the Discriminator, which are trained together in a game-theoretic framework. The two networks are implemented as feedforward neural networks. We can generate the synthetic dataset from the real datasets with good outcomes and bad outcomes and four time steps, or three time breaks.

The *Generator* takes a random noise vector as input and produces synthetic data. It consists of three fully connected layers with ReLU activation functions. The first layer is a fully connected (dense) layer that takes an input vector of size 100 and produces an output vector of size 128 (here using the torch.nn library: `nn.Linear(100, 128)`). After each linear layer, a ReLU activation function is applied. The second layer is `nn.Linear(128, 256)`, that takes the 128-dimensional output from the previous layer and produces a 256-dimensional output. The third layer is `nn.Linear(256, 4)`: This is the final fully connected layer of the Generator. It takes the 256-dimensional input and produces a 4-dimensional output to match the dimension of our output data (4 time steps).

The *Discriminator* takes data as input and outputs a probability score indicating whether the input is real or fake. We define the Discriminator as three fully connected layers with ReLU activation functions and a sigmoid activation in the final layer to produce a probability score. The first layer is a fully connected `nn.Linear(4, 128)` that takes an input vector of size 4 (for the 4 time steps), and produces an output vector of size 128. A ReLU activation is applied. The second layer is `nn.Linear(128, 64)`, a ReLU function is applied, and the last layer is `nn.Linear(64, 1)` with a sigmoid activation function for the probability of the input being real or fake.

We set standard hyperparameters for GAN structures: a learning rate of 0.0002, a batch size of 64 and 10.000 epochs. Both networks are set to use a Binary Cross-Entropy loss function and an Adam optimizer. The training loop alternates between training the Discriminator and training the Generator to ensure that the skill levels are similar. In each iteration, a random noise is generated as input for the Generator, fake data is generated by passing the noise through the Generator. Real patient data and fake data are used to train the Discriminator to correctly classify real as 1 and fake as 0. The Generator is trained to produce data that the Discriminator classifies as real (aiming for a loss close to 1). The “equilibrium” point in this min-max game occurs when the Generator produces data that is indistinguishable from real data, and the Discriminator assigns a probability of 0.5 to both real and generated data. We then generate 100 fake patients using the Generator trained previously and plot the results for each cohort with standard noise. A subset of 20 patients is illustrated in the figure below. Good outcome patients are also illustrated.

6.2 Modelling Atlases Using the MCMC-SAEM Algorithm

In this section, we propose to model the two population trajectories based on Mantoux et al’s Mantoux (2022a) adaptation of mixed-effect models for longitudinal datasets of unknown time breaks.

We define a hierarchical structure which decouples the average trajectory of the population from the individual-level variability Mantoux et al. (2021). In particular, we are interested in the structural breaks within the population dynamics that could correspond to either a treatment response or an infra-clinical disease progression. The trajectory of each individual is modeled as a piecewise affine trajectory.

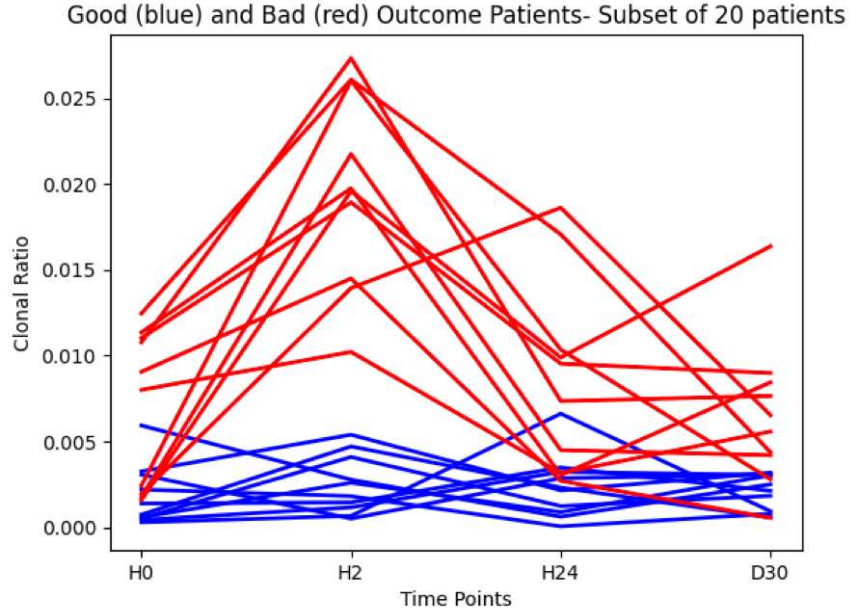


Figure 15: GAN generated good and bad outcome patients.

Our data is summarized by y and t . Here, we define y as the collection of all observations of each individual, i.e. y is a vector of 4 ctDNA ratios (or values) for each patient. We define t as the collection of observation times, which are H0, H2, H24 and D30 for the plasma samples. Fixed and random effects are not obvious when plotting both individual trajectories. The framework was previously adapted by Clément Mantoux Mantoux (2022b) Mantoux et al. (2022). The model first defines an average population trajectory $D(t)$. This trajectory is taken as a d -dimensional continuous piecewise linear curve, which allows handling structural breaks in the disease progression or changes in a treatment. In our case, the dimension is 1 and the collection of observation times is time 0, 1, 2 and 3, where the thermal ablation was performed between time 0 and 1 before the theoretical release of ctDNA.

The framework is particularly adapted to the clinical setting where the structural breaks (thermal ablation disease recurrence) are not known.

This work is based on models published by Schiratti et al SCHIRATTI et al. (2015). In this article, the authors proposed a Bayesian mixed-effects model designed for analyzing longitudinal manifold-valued data. Their model allows for the estimation of a group-average trajectory in the measurement space – the trajectory of ctDNA for good or bad outcomes patients in our case - using Riemannian geometry tools. Using the tools of geometry allow us in this case to make no assumption about the trajectories of ctDNA. Most importantly, individual trajectories derive from the group average and are defined as a parallel to the group average. A stochastic version of the EM (the SAEM) is used to estimate the parameters of the model.

Building on Juliette Chevallier’s thesis Chevallier (2019), adaptation is applied as follows: the average population trajectory is parameterized by a reference point $(p_0, t_{B,0})$ in \mathbb{R} and K breaks with time breakpoints t_1, t_2, \dots, t_{K-1} . Each breakpoint has a corresponding position p_i on the

```

Initialize  $\theta_0, z_0$  and  $S_0$ 
repeat
  Sample  $z_{t+1} \sim q(\cdot | z_t; \theta_t)$  from the MCMC kernel  $q(z | z', \theta_t)$  targeting  $p(z | y, \theta_t)$ 
  Update  $S_{t+1} = (1 - \gamma_t)S_t + \gamma_t S(y, z_{t+1})$ 
  Find  $\theta_{t+1} \in \operatorname{argmax}_{\theta} \langle S_{t+1}, \Phi(\theta) \rangle + \Psi(\theta)$ 
until convergence
return  $\theta_T, (z_t)_{t=1}^T$ 

```

Figure 16: MCMC-SAEM algorithm applied to ctDNA, adapted from Juliette Chevallier

trajectory $D(t_i)$. The measurements of ctDNA at different points are modeled as noised observations of the hidden D_i trajectory. Each trajectory is defined by both a space shift and a time reparameterization of the average population trajectory $D(t)$. Each trajectory can be defined by its hidden latent variables, i.e. the parameters.

A population trajectory is fully defined by the list of break times and the list of trajectory values at each break. The trajectory is hypothesized to be affine in between the breaks. These define the first three parameters defined as $p0_bar$, the list of trajectory values at the first break, $t0_bar$ is the list of break times and $v0_bar$ is the list of slopes for each coordinate and each piece of the trajectory. We will then add a noise on the population trajectory. These population parameters are stacked in the variable (x, y, θ)

Each individual is defined by how its trajectory deviates from the average ($p0, t0, v0$) with the following hidden parameters: τ gives a temporal shift between individual i and the population trajectory, $\epsilon [i,j]$ represents the logarithmic acceleration factor of individual i on the piece j of the trajectory, ϕ is the spatial shift between the population trajectory and the individual trajectory. These individual parameters are stacked in the variable z .

No missing data were imputed in this section. The model parameters are computed using the Maximum A posteriori estimator (MAP), which was computed with the MCMC-SAEM algorithm.

For feasibility purposes, we use the previously generated synthetic data as if we had enrolled 200 patients, with simple labels as ‘progressive’ (bad outcome) and ‘non-progressive’ (good outcome).

We aim to later test ctDNA in real-life settings through a multicentric clinical trial.

We apply the MCMC-SAEM and define the number of SAEM steps for the parameter estimation : 1000000 ; the number of MCMC iterations to estimate the posterior mean of $(z|y)$: 10000 ; the number of MCMC iterations used in the marginal likelihood computation : 100000. The previous work proposes to initialize the MCMC-SAEM by taking a linear regression of each feature across time, and to use a Symmetric Random Walk Metropolis Hastings within Gibbs sampler with Gaussian transitions. The variance of the proposal Gaussian transitions are tuned along the SAEM steps to reach a desired Metropolis acceptance rate. The complete procedure for the marginal likelihood estimation was summarized:

In the course of our clinical study, we embarked on a dual track of research aimed at both helping us select patients prior to treatment and helping us better understand the biological effects of thermal ablation. Circulating tumor DNA proved to be effective as a predictive biomarker prior to the thermal ablation. Unfortunately, the four time points evaluated were insufficient for a thorough understanding of the dynamics.

We hypothesized a pronounced increase in ctDNA levels following treatment due to cell destruction through necrosis, apoptosis, and exosome formation, along with the triggered immunity reaction. We hypothesized that patients who would not respond in the long term were patients with undiagnosed untreated disease, i.e. who would have a steep fall in ctDNA before a steep uptake. In parallel, good outcome patients would have sustainability in the minimal residual ctDNA found

```

Run a first Metropolis within Gibbs chain  $z_{i,1}^{(0)}, \dots, z_{i,M_0}^{(0)}$  targeting the distribution
of  $p(z_i | y_i, \theta)$ 
Compute the empirical mean  $\mu_{IS}$  and covariance  $\Sigma_{IS}$  of  $(z_i^{(t)})_{1 \leq t \leq M_0}$ 
Define the importance distribution  $q_{IS,i} = \mathcal{N}(\mu_{IS}, \Sigma_{IS})$ 
Run a second Metropolis within Gibbs chain to draw  $M_1$  samples  $z_{i,1}^{(1)}, \dots, z_{i,M_1}^{(1)}$ 
targeting  $p(z_i | y_i, \theta)$ 
Draw  $M_2$  samples  $z_{i,1}^{(2)}, \dots, z_{i,M_2}^{(2)}$  from  $q_{IS,i}$ 
Initialize  $\hat{p}_0 = 0$ 
repeat
  Define  $\hat{h}_t$ , the approximation of  $h$  obtained by replacing  $p(y_i | \theta)$  with  $\hat{p}_t$  in Equation
  (7.2)
  Compute  $\hat{p}_{t+1}$  using Equation (7.1) by replacing  $h$  with  $\hat{h}_t$ 
until convergence
Compute the RMSE estimator for  $\hat{p}_{\text{final}}$ 
return  $\hat{p}_{\text{final}}$  and the RMSE estimator

```

Figure 17: From Mantoux et al, bridge sampling procedure for the marginal likelihood estimation.




Figure 18: Outcomes and trajectories using the SAEM method.

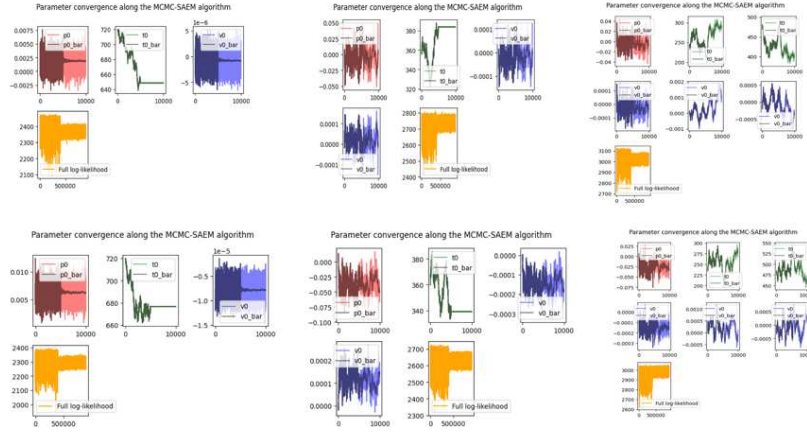


Figure 19: Parameters for $K=0, 1$ and 2 for good outcome (first line) and bad outcome (second line) patients.

throughout time.

Both hypothesized trajectories were produced by the MCMC-SAEM algorithm, with, however, some specifics. Trajectories aligned with our expectations, with a sharp increase observed among patients with unfavorable outcomes and a steady maintenance of low levels among responders, and even showed the small uptake following thermal ablation. Additionally, the detection of temporal breaks, particularly around the 300/400 timepoints, presents an intriguing result for further investigation into the dynamics of untreated cancer progression. This timebreak, at approximately two weeks after treatment, suggests a change in dynamics that is concordant with known timely evolutions in colorectal cancer. Two weeks is typically the intercurse delay between chemotherapy and targeted therapies lines, and is also the delay for interrupting treatments in the peri operative setting. Whether this result is artefactual or not could be the subject of a prospective study on a broader cohort of patients treated with thermal ablation.

By marking a break at Hour 400, the model provides important insight into what we could expect from tumor DNA dynamics in regards to the underlying biological phenomena induced by thermal ablation.

6.3 Discussion

Using a GAN on 20 points in 4 dimensions seems like a rather complicated approach. GANs are known to be unstable during the training process, stemming from the adversarial nature of the training where both networks compete against each other. Since such networks are usually trained on massive data, our GAN might fail to converge towards a plausible population trajectory for both good outcome and bad outcome patients. It may mimic the statistical properties of our small training data based on the clinical study conducted for this work, without capturing the complex biological relationships between ctDNA, thermal ablation and disease recurrence. Recent research is specifically aimed at the mode collapse issue with GANs, illustrated by examples such as the Dynamic GAN framework proposed by Luo et al²³. In this publication, they propose to address the diversity of the generated samples (which is often lower than realistic), by detecting collapsed

samples and training a dynamic model in parallel to the competition. These new frameworks could be assessed on bioclinical data such as ctDNA. Simpler models could be explored, such as estimating a multivariate Gaussian of dimension 4 for each class (good/bad outcome), and could be compared with the GAN results. Nevertheless, this enabled us to have a virtual population to try the SAEM model.

The model proposed by Juliette Chevallier and Clement Mantoux weren't tested in this specific context: on average, there were 15 points in each individual's trajectory. It's uncertain if it works with only 4 observations, especially when two of them (H0 and H2) are very close together in time. It wouldn't be surprising if the model treats the measurement gap between these two points as noise. This could pose a significant obstacle to applying the model as-is in this situation. Adaptation might be necessary (e.g., removing noise in observations by imposing a very small epsilon sigma, which should be feasible via the prior on parameters), and/or considering how to detect slope changes if the time scale changes (e.g., switching to a logarithmic time scale).

The mixed effect model might seem overly complex for the intended application, but we could use it in a prospective study with more than 4 samples per patient. It is difficult to determine if breakpoints occur at a specific point in the interval, for example, halfway, or at point H24. Placing many breakpoints in this interval indicates that the model might be somewhat lost, especially since the trajectory segments between two breakpoints in the interval [H24, D30] don't seem to serve a clear purpose. Additional information would be needed to extrapolate what happens in the middle of the interval.

A suggestion is to create a simple mixed-effects model where the breakpoint dates are known and fixed (namely, H0, H2, H24, D30). This should be easier to estimate than the current model and doesn't assume breakpoints between two points, which is difficult to verify. Comparing 4 models: no breakpoint, breakpoint at H2, breakpoint at H24, breakpoint at H2 and H24 could provide valuable insights.

In conclusion, we identify patients with bad outcome as patients who have a ascending ctDNA after thermal ablation, and patients with good outcome as a steady low-value ctDNA throughout follow-up. This could be used in clinical practice, and should be further investigated in large prospective studies.

7 REFERENCES

References

- Kustanovich A, Schwartz R, Peretz T, and Grinshpun A. Life and death of circulating cell-free dna. *Cancer biology therapy*, 20(8), 2019. ISSN 1555-8576. doi: 10.1080/15384047.2019.1598759. URL <https://pubmed.ncbi.nlm.nih.gov/30990132/>. publisher: Cancer Biol Ther PMID: 30990132.
- François-Clément Bidard, Nicolas Kiavue, Marc Ychou, Luc Cabel, Marc-Henri Stern, Jordan Madic, Adrien Saliou, Aurore Rampanou, Charles Decraene, Olivier Bouché, Michel Rivoire, François Ghiringhelli, Eric Francois, Rosine Guimbaud, Laurent Mineur, Faiza Khemissa-Akouz, Thibault Mazard, Driffa Moussata, Charlotte Proudhon, Jean-Yves Pierga, Trevor Stanbury, Simon Thézenas, and Pascale Mariani. Circulating tumor cells and circulating tumor dna detection in potentially resectable metastatic colorectal cancer: A prospective ancillary study to the uni-cancer prodige-14 trial. *Cells*, 8(6):516, 5 2019. ISSN 2073-4409. doi: 10.3390/cells8060516. PMID: 31142037 PMCID: PMC6627974.
- Anders K. Boysen, Niels Pallisgaard, Christina S. A. Andersen, and Karen-Lise G. Spindler. Circulating tumor dna as a marker of minimal residual disease following local treatment of metastases from colorectal cancer. *Acta Oncologica*, 59(12):1424–1429, 12 2020. ISSN 0284-186X. doi: 10.1080/0284186X.2020.1806357. publisher: Taylor Francis *eprint* : <https://doi.org/10.1080/0284186X.2020.1806357>.
- A. Cervantes, R. Adam, S. Roselló, D. Arnold, N. Normanno, J. Taïeb, J. Seligmann, T. De Baere, P. Osterlund, T. Yoshino, E. Martinelli, and ESMO Guidelines Committee. Electronic address: clinicalguidelines@esmo.org. Metastatic colorectal cancer: Esmo clinical practice guideline for diagnosis, treatment and follow-up. *Annals of Oncology: Official Journal of the European Society for Medical Oncology*, 34(1):10–32, 1 2023. ISSN 1569-8041. doi: 10.1016/j.annonc.2022.10.003. PMID: 36307056.
- Chao Cheng, Motoko Omura-Minamisawa, Yun Kang, Takamitsu Hara, Izumi Koike, and Tomio Inoue. Quantification of circulating cell-free dna in the plasma of cancer patients during radiation therapy. *Cancer Science*, 100(2):303–309, 2 2009. ISSN 1349-7006. doi: 10.1111/j.1349-7006.2008.01021.x. PMID: 19200259.
- Juliette Chevallier. *Statistical models and stochastic algorithms for the analysis of longitudinal Riemannian manifold valued data with multiple dynamic*. PhD thesis, 9 2019. URL <https://www.theses.fr/2019SACLX059>. [Online; accessed 2024-03-05].
- Brett L. Ecker, Jasme Lee, Lily V. Saadat, Thomas Aparicio, Florian E. Buisman, Vinod P. Balachandran, Jeffrey A. Drebin, Kiyoshi Hasegawa, William R. Jarnagin, Nancy E. Kemeny, T. Peter Kingham, Bas Groot Koerkamp, Norihiro Kokudo, Yutaka Matsuyama, Guillaume Portier, Leonard B. Saltz, Kevin C. Soares, Alice C. Wei, Mithat Gonen, and Michael I. D’Angelica. Recurrence-free survival versus overall survival as a primary endpoint for studies of resected colorectal liver metastasis: a retrospective study and meta-analysis. *The Lancet Oncology*, 23(10):1332–1342, 10 2022. ISSN 1470-2045, 1474-5488. doi: 10.1016/S1470-2045(22)00506-X. publisher: Elsevier PMID: 36058227.

- Matthew Loft, Yat Hang To, Peter Gibbs, and Jeanne Tie. Clinical application of circulating tumour dna in colorectal cancer. *The Lancet Gastroenterology Hepatology*, 8(9):837–852, 9 2023. ISSN 2468-1253, 2468-1156. doi: 10.1016/S2468-1253(23)00146-2. publisher: Elsevier PMID: 37499673.
- P. Mandel and P. Metais. [nuclear acids in human blood plasma]. *Comptes Rendus Des Seances De La Societe De Biologie Et De Ses Filiales*, 142(3-4):241–243, 2 1948. ISSN 0037-9026. PMID: 18875018.
- Clément Mantoux. *Modélisation statistique et inférence pour les populations de réseaux de connectivité cérébrale et les données longitudinales*. PhD thesis, 9 2022a. URL <https://www.theses.fr/2022IPPAX072>. [Online; accessed 2024-03-05].
- Clément Mantoux. *Statistical Modeling and Inference for Populations of Networks and Longitudinal Data*. PhD thesis, 9 2022b. URL <https://theses.hal.science/tel-04105512>. [Online; accessed 2024-02-28].
- Clément Mantoux, Baptiste Couvy-Duchesne, Federica Cacciamani, Stéphane Epelbaum, Stanley Durrleman, and Stéphanie Allasonnière. Understanding the variability in graph data sets through statistical modeling on the stiefel manifold. *Entropy (Basel, Switzerland)*, 23(4):490, 4 2021. ISSN 1099-4300. doi: 10.3390/e23040490. PMID: 33924060 PMCID: PMC8074266.
- Clément Mantoux, Stanley Durrleman, and Stéphanie Allasonnière. Asymptotic analysis of a matrix latent decomposition model. *ESAIM: Probability and Statistics*, 26:208–242, 2022. doi: 10.1051/ps/2022004. publisher: EDP Sciences.
- Emilie Moati, Valérie Taly, Audrey Didelot, Géraldine Perkins, Hélène Blons, Julien Taieb, Pierre Laurent-Puig, and Aziz Zaanani. Role of circulating tumor dna in the management of patients with colorectal cancer. *Clinics and Research in Hepatology and Gastroenterology*, 42(5):396–402, 10 2018. ISSN 2210-741X. doi: 10.1016/j.clinre.2018.03.002. PMID: 29627453.
- Raja R. Narayan, Debra A. Goldman, Mithat Gonen, Jonathan Reichel, Kety H. Huberman, Sandeep Raj, Agnes Viale, Nancy E. Kemeny, Peter J. Allen, Vinod P. Balachandran, Michael I. D’Angelica, Ronald P. DeMatteo, Jeffrey A. Drebin, William R. Jarnagin, and T. Peter Kingham. Peripheral circulating tumor dna detection predicts poor outcomes after liver resection for metastatic colorectal cancer. *Annals of Surgical Oncology*, 26(6):1824–1832, 6 2019. ISSN 1534-4681. doi: 10.1245/s10434-019-07201-5. PMID: 30706231 PMCID: PMC6511310.
- Robbert S. Puijk, Alette H. Ruars, Laurien G. P. H. Vroomen, Aukje A. J. M. van Tilborg, Hester J. Scheffer, Karin Nielsen, Marcus C. de Jong, Jan J. J. de Vries, Babs M. Zonderhuis, Hasan H. Eker, Geert Kazemier, Henk Verheul, Bram B. van der Meijs, Laura van Dam, Natasha Sorgedrager, Veerle M. H. Coupé, Petrousjka M. P. van den Tol, Martijn R. Meijerink, and COLLISION Trial Group. Colorectal liver metastases: surgery versus thermal ablation (collision) - a phase iii single-blind prospective randomized controlled trial. *BMC cancer*, 18(1):821, 8 2018. ISSN 1471-2407. doi: 10.1186/s12885-018-4716-8. PMID: 30111304 PMCID: PMC6094448.
- Nicolas Pécuchet, Eleonora Zonta, Audrey Didelot, Pierre Combe, Constance Thibault, Laure Gibault, Camille Lours, Yves Rozenholc, Valérie Taly, Pierre Laurent-Puig, Hélène Blons, and Elizabeth Fabre. Base-position error rate analysis of next-generation sequencing applied to circulating tumor dna in non-small cell lung cancer: A prospective study. *PLoS Medicine*, 13(12):e1002199, 12 2016. ISSN 1549-1277. doi: 10.1371/journal.pmed.1002199. PMID: 28027313 PMCID: PMC5189949.

Jean-Baptiste SCHIRATTI, Stéphanie ALLASSONNIERE, Olivier Colliot, and Stanley DURRLEMAN. Learning spatiotemporal trajectories from manifold-valued longitudinal data. volume 28. Curran Associates, Inc., 2015. URL https://papers.nips.cc/paper_files/paper/2015/hash/186a157b2992e7daed3677ce8e9fe40f-Abstract.html. [Online; accessed 2024-01-22].

Jeanne Tie, Yuxuan Wang, Joshua Cohen, Lu Li, Wei Hong, Michael Christie, Hui Li Wong, Suzanne Kosmider, Rachel Wong, Benjamin Thomson, Julian Choi, Adrian Fox, Kathryn Field, Matthew Burge, Jenny Shannon, Dusan Kotasek, Niall C. Tebbutt, Christos Karapetis, Craig Underhill, Andrew Haydon, Joy Schaeffer, Janine Ptak, Cristian Tomasetti, Nicholas Papadopoulos, Kenneth W. Kinzler, Bert Vogelstein, and Peter Gibbs. Circulating tumor dna dynamics and recurrence risk in patients undergoing curative intent resection of colorectal cancer liver metastases: A prospective cohort study. *PLoS medicine*, 18(5):e1003620, 5 2021. ISSN 1549-1676. doi: 10.1371/journal.pmed.1003620. PMID: 33939694 PMCID: PMC8128260.

Jeanne Tie, Joshua D. Cohen, Kamel Lahouel, Serigne N. Lo, Yuxuan Wang, Suzanne Kosmider, Rachel Wong, Jeremy Shapiro, Margaret Lee, Sam Harris, Adnan Khattak, Matthew Burge, Marion Harris, James Lynam, Louise Nott, Fiona Day, Theresa Hayes, Sue-Anne McLachlan, Belinda Lee, Janine Ptak, Natalie Silliman, Lisa Dobbyn, Maria Popoli, Ralph Hruban, Anne Marie Lennon, Nicholas Papadopoulos, Kenneth W. Kinzler, Bert Vogelstein, Cristian Tomasetti, and Peter Gibbs. Circulating tumor dna analysis guiding adjuvant therapy in stage ii colon cancer. *New England Journal of Medicine*, 386(24):2261–2272, 6 2022. ISSN 0028-4793. doi: 10.1056/NEJMoa2200075. publisher: Massachusetts Medical Society *eprint* : <https://doi.org/10.1056/NEJMoa2200075> PMID : 35657320.

Part Four: Pre-requisites To Automatic Endovascular Navigation Systems: Device Segmentation And Curvature-Preserving Arborescent Representation Of Vascular Systems

August 15, 2024

Contents

1	Navigating within the patient's vascular network	2
1.1	Introduction to endovascular navigation	2
1.2	Material and method	4
1.2.1	Patient selection	5
1.2.2	Deep learning model	5
1.3	Evaluation	9
2	Curvature-preserving arborescent representation of cerebral vasculature	12
2.1	Introduction	12
2.2	Pre-processing of the per operative 3D acquisition	13
2.3	Graphical representation of the vascular network	14
2.4	Planar graphical representation of the vascular network	18
3	Extension towards the origins of the vascular network: live embryology	21
3.1	Introduction	21
3.2	Dataset and study objectives	22




Figure 1: PhD map.

1 Navigating within the patient’s vascular network

This section refers to the *purple projects*.

This section was adapted from the following publications:

How Artificial Intelligence will Reshape our Interventional Units published in Cardiovasc Intervent Radiol (2023). Boeken et al. (2023a)

Automated catheter segmentation for endovascular radioscopy images with deep [morphological] networks- To be submitted

Curvature-preserving arborescent representation of cerebral vasculature- To be submitted.

Special thanks for this section: for the ongoing collaborative work with Jean Feydy, Alisa Kugusheva, Julien Nguyen Van, Philippine Cordelle.

1.1 Introduction to endovascular navigation

After patient selection, our work focused on the per-interventional data.

The first reported angiography of blood vessels was in 1896, in Vienna, performed on the hand of a cadaver. The first *cerebral* angiography was performed in 1927, by Moniz. This was the only diagnostic technique that allowed the investigation of vascular disorders until the computed tomography was introduced in 1975. The pioneer work enabled to visualize cerebral vessels in live patients, creating an “Arterial Encephalography”. The method of obtaining cerebral angiograms has evolved over the years. In the 1960s, the procedure involved puncturing the carotid artery and injecting a contrast medium. The advent of Digital Subtraction Angiography (DSA) in the 1980s marked a major advancement in the field, and more recently, the use of 3D angiography has become increasingly prevalent.

This technique allows ultra-high-resolution vascular images that have not yet been used to navigate autonomously. Being able to interpret those live images is a prerequisite for automation and robotic navigation.

Endovascular image-guided interventions have revolutionized the medical field, offering targeted treatments within organs through vessel navigation. This minimally invasive approach relies on live




Figure 10: First "arteriogram" of an amputated band (Hascbek and Lindenthal 1896).

Figure 2: 1927 arteriogram





Figure 3: Contemporary Unit

fluoroscopic imaging and radio-opaque devices to guide physicians during procedures. Today, most procedures are performed by a *manual* operator who inserts these devices while analyzing the live images, using a two-degree of freedom action: the push-pull action and the torque action.

For example, uterine artery embolization, a well-established interventional radiology treatment, has consistently demonstrated its efficacy in treating women with symptomatic fibroids Barat et al. (2024). The procedure involves the precise injection of embolization particles directly into the uterine arteries, navigating catheters through a single artery puncture site.

Robotic assistance holds the potential to change the manual aspect of these therapies Boeken et al. (2023a). Due to the apparent simplicity of the fundamental actions involved (such as pushing and twisting), endovascular approaches are particularly prone to shift towards teleoperated interventions Crinnion et al. (2021). The technology promises to enhance precision, reduce radiation exposure, and allow remote navigation within the vascular network. It is already being used in complex and high-risk situations, such as the treatment brain aneurysms or carotid artery stenting, where precise navigation and stability are crucial.

These solutions may benefit from recent advances in Artificial Intelligence and Computer Vision, evolving from robotic assistance towards more autonomy. The concept of autonomous interventions aims at developing systems that can analyze live images during the intervention and enable some degree of feedback loop or decision process. Computer vision models have undergone significant adaptation to cater to the immense interest in Diagnostic Imaging. In the realm of Interventional Imaging, various AI-based models are readily available for live image analysis, offering valuable insights like predicting vascular flow dynamics and minimizing misregistration artifacts Najafi et al. (2023) Boeken et al. (2023b). A significant challenge that Computer Vision tackles is precisely identifying guidewires, catheters, and contrast positions, effectively narrowing the gap between robotic assistance and the interpretation of live images.

These robots are unable to integrate the live image at the core of image-guided therapies. Today, they mostly integrate the *manual* aspect of the therapy without the *visual perception* guiding the treatment. In this sense, such robots are similar to well-established surgical solutions for open or laparoscopic approaches. Computer vision features such as automatic catheter detection could be implemented in contemporary robotic solutions. As an example, one robotic software feature named ‘Rotate and Retract’ is an automated robotic movement that was FDA approved. Other automated actions are currently being evaluated, including the ‘Spin’, ‘Wiggle’, ‘Dotter’, and ‘Constant speed’ features.

Catheter tip detection and segmentation are two prerequisites to such live visual interpretation. Previous solutions were recently proposed, including a temporary transformer network for guidewire segmentation Zhang et al. (2021) and a tip detection model for cerebral angiography Ghosh et al. (2023). Pelvic interventions, such as uterine artery embolization, could be an area of interest for automation, in the context of remote treatment for patients from regions without interventional radiologists or in the context of acute hemorrhage. Most published studies were performed on either neurovascular or cardiovascular interventions Zhang et al. (2022).

In the present study, we propose to train and evaluate a deep morphological network for live catheter segmentation in uterine artery embolization.

1.2 Material and method

1.2.1 Patient selection

This study was IRB approved (IRB # #00011928). All consecutive patients who underwent uterine artery embolization (UAE) between January 2022 and December 2022 were screened.

Inclusion criteria comprised all women who underwent UAE at our tertiary center and patients were excluded if they refused to participate or if they were under 18 or under guardianship.

All procedures were performed by experienced radiologists from the interventional radiology department and live radioscopic images were prospectively recorded. Standard parameters were used without specific treatment changes due to the study. UAE was either performed for post-partum hemorrhage, adenomyosis or symptomatic uterine fibroids. Post-procedure, patients were hospitalized for pain management and discharged the following day with painkiller prescriptions. All DICOM data were fully anonymized and images were retrieved and annotated by three independent investigators using the GIMP software.

A total of 152 patients were included. For each patient, a randomly selected single image stemming from the scopic video was retrieved containing the tip of the catheter and/or the micro-catheter at the injection point. The full was included for further analysis. Individual images were extracted from the DICOM metadata and subsequently converted into standard grayscale files to facilitate analysis and processing. Grayscale images contain pixel values ranging from 0 (black) to 255 (white), representing the intensity of the image at each pixel location. The images vary in sizes from 512x512 to 1024x1024.

The dataset was divided in a training set (period 1) of 112 images, a validation set of 20 images (period 2) and an independent test dataset of 20 images (period 3) for performance analysis.

The corresponding segmentation masks were manually annotated for each image and represent a binary image with background pixels as black (0) and a catheter/microcatheter mask as white (1). Due to the limitation of data, and in order to avoid overfitting and to make the model more robust, we performed several data augmentation techniques. The transforms included: random horizontal flip, random perspective, random rotation, random brightness contrast, Gaussian noise, and a center crop. The applied data augmentation algorithms provided more diversity to the datasets, giving 448 images in the training set and 80 images in the validation set. The result of data augmentation is represented in the figure below. An example is provided in the figure below.

1.2.2 Deep learning model

The model used is a U-Net with grouped convolutions Cohen and Welling (2016). The model used is a U-Net specifically adapted to annotated medical image segmentation tasks Ronneberger et al. (2015). Briefly, the U-Net model architecture starts with an encoder part whose purpose is to downscale the images while increasing the number of feature maps. The downscaling step is repeated 4 times in total to reach the bottleneck. The second part, also called the decoder, is used to upscale the images while decreasing the number of feature maps in a mirroring way compared to the encoder. Skip connections between the similarly shaped feature maps of the encoder and decoder make it possible for the network to pass useful information for performant upscaling.

As it was previously shown, the vanilla U-Net performs poorly on the catheter segmentation task. Using grouped convolution instead of the classical one significantly increases the performance. In this study, we show results for a model with grouped convolutions.

The mask retrieved in the output is compared to the ground truth mask with several loss functions. Because the pixels involved in the segmentation task of the catheter are very rare




Figure 4: Examples of transformations – Up :images, Down: corresponding mask - From left to right : Original image, Random perspective, Random rotation, Horizontal Flip, Vertical Flip




Figure 5: Examples of transformations – Up :images, Down: corresponding mask - From left to right : Original image, Random perspective, Random rotation, Horizontal Flip, Vertical Flip




Figure 6: More transformations

Raw data



Mask



Example of flip-rotation transformation for data augmentation




Figure 7: More transformations

compared to the overall background pixels, such tasks are considered unbalanced problems. Specific loss functions were proposed to address such class imbalance related issues, such as weighted cross-entropy, Dice loss or Sensitivity-Specificity Sudre et al. (2017). All these functions use a binary classification (catheter / background). The Dice loss function is based on the Sorensen-Dice Dice (1945) Sørensen (1948) coefficient that measures the overlap between the prediction and the ground truth, and was later developed as a generalized Dice loss function for multiple class segmentation. The Dice was performed on $A = \text{catheter}$ and $B = \text{background}$.

More recently, the centerline Dice, named shortly clDice loss function was introduced to preserve the topology of tubular structures such as arteries and catheters Shit et al. (2021). In this paper, Shit et al propose to assess the intersection between the mask and its morphological skeletal representation and demonstrate topology preservation for 2D and 3D segmentation. While the classical Dice score treats all pixels equally, in the task of catheter segmentation the connected components on the image are subject to a specific focus in clDice score. We also train with the active contour with elastica (ACE) loss function that incorporates elastica (curvature and length) and region information as geometrically-natural constraints for the image segmentation tasks Chen et al. (2020) .

Backpropagation is then performed across the network to update its weights. We used the Adam optimizer and started with a learning rate of 0.01. The learning rate was gradually decreased, being halved every 10 epochs. In total the model was trained on no more than 30 epochs as a plateau was reached for the value of the loss function.

$$\text{Dice Score} = \frac{2 \times |A \cap B|}{|A| + |B|}$$

where A and B are two distinct sets. ($A = \text{catheter}$ and $B = \text{background}$).

For the clDICE score, we consider two binary masks: the ground truth mask (V_L) and the predicted segmentation masks (V_P) according to Shit et al. First, the skeletons S_P and S_L are extracted from V_P and V_L respectively. Then the fraction of S_P that lies within V_L , which they call *Topology Precision* or $\text{Tprec}(S_P, V_L)$, and *Topology Sensitivity* or $\text{Tsens}(S_L, V_P)$ are defined below;

$$\text{Tprec}(S_P, V_L) = \frac{|S_P \cap V_L|}{|S_P|}; \quad \text{Tsens}(S_L, V_P) = \frac{|S_L \cap V_P|}{|S_L|} \quad (1)$$

$$\text{clDice}(V_P, V_L) = 2 \times \frac{\text{Tprec}(S_P, V_L) \times \text{Tsens}(S_L, V_P)}{\text{Tprec}(S_P, V_L) + \text{Tsens}(S_L, V_P)} \quad (2)$$

The model was trained on around 30 epochs (see the figure below) .

The model was trained on Tesla T4 GPUs. The images were center cropped to the size of 480x480 and their pixel values ranged from 0 to 1. The training was performed with a batch size of 4.

1.3 Evaluation

We perform the evaluation of our model on a separate test set consisting of 20 images. We use IoU (intersection over union) and clDice score as an evaluation metric. The results of the segmentation are represented in the table below. This section is still under investigation. ***Performances are provided for IoU (intersection over union), and clDICE are coming next as this is an ongoing project.***




Figure 8: Loss using cIDICE

Loss function	IoU on the test set (\uparrow)	cIDICE
Dice	0.00123	
BCE	0.00172	
BCE + Dice	0.0016	
BCE + cIDice	0.00109	
Dice + cIDice	0.0017	

Figure 9: Performance according to loss functions. **Note that cIDICE are still under investigation.**



Figure 10: Target / prediction on one patient.




Figure 2 Devices are labeled and tracked by artificial intelligence assistance software, Neuro-Vascular Assist. (A) A filter wire tip is indicated by a triangle. (B) A filter is indicated by a crosshair. (C) A guiding catheter tip is indicated by a square.

2

Sakakura Y, et al. *J NeuroIntervent Surg* 2024;0:1–5. doi:10.1136/jnis-2024-021600

Figure 11: Neuro-Vascular Assist (iMed technologies, Tokyo, Japan), from Sakakura et al.

Note that Sakakura et al published the first case of carotid artery stenting using their proprietary software tracking wires, guiding catheters, and embolic protection devices Sakakura et al. (2024). The full segmentation of the device was not needed for this study, a bounding box was sufficient to assist neuro-interventional radiologists in their procedure.

While this study is ongoing, understanding live images is paramount to automation. This will bridge the gap between pre-therapeutic planification (notably from 3D images) to robotic assistance (with motorized arms).

2 Curvature-preserving arborescent representation of cerebral vasculature

2.1 Introduction

Motivation. Current imaging techniques for cerebral vasculature include mini-invasive high-resolution selective three-dimensional subtracted angiographies. These three-dimensional contrast-enhanced images are obtained *prior* to the treatment of several vascular disorders such as aneurysms, malformations, fistulas, strokes and stenosis. They provide essential data for treatment planning, especially regarding the vascular access route towards the target disorder.

Endovascular navigation is based on the insertion of several coaxial devices including guides, micro-guides, sheaths, catheters and microcatheters. All of these devices present specific torque, support and pushability characteristics that will determine the difficulties in accessing target points. These challenges are inherently correlated to the curvature and length of the vessel route, and the parietal friction along the way.

Specifications. Unfortunately, 3D rendering of the patient anatomy does not provide an understandable representation of possible endovascular routes and related challenges. This is the reason why we aim at providing a simple planar arborescent representation of the full cerebral vasculature that can help interventional radiologists to assess obstacles prior to endovascular treatment, using a single snapshot of the patient anatomy. Our work is also motivated by pioneering anatomical studies of the brain vasculature (Kahilogullari et al., 2012), performed via physical dissections of dozens of human brains, that we would now like to automate in a mini-invasive way.

Starting from a 3D angiography, we intend to create a planar representation of the 3D vascular structure that:

- Preserves the real length and diameter of the vessel segments.
- Approximates the angles and take-offs of the artery branches.
- Estimates the tortuosity of the artery segments through curvature.
- Estimates the accumulated torque of the device along the route.
- Optimizes the utilization of the available representation 2D-space efficiently.
- Ensures that there is no overlap in the two-dimensional (2D) representation of the arteries.

We note that the creation of blood vessel maps is an active field of research, and refer to Eulzer et al. (2022) for a modern review. **What sets our work apart is the attention paid to the vessel lengths and curvatures, which are often discarded by other methods but play a crucial role in endovascular interventions.**

Cohort. Our study population consists of fifteen consecutive patients followed for unruptured intracranial aneurysms in a single tertiary center. The procedures were performed with the patient under general anesthesia. The first step in the process involves positioning the C-arm angiography system around the patient’s head. This equipment consists of a C-shaped arm with an X-ray source and detector that can rotate around the patient to capture images from multiple angles




Figure 12: Real vascular anatomy from cadaveric dissection. This resembles the current study’s objectives.

(Cone Beam CT). The first cone-beam computed tomography (CBCT) acquisition is performed without the injection of contrast material. This initial acquisition serves to create a baseline or “mask” image of the patient’s head and intracranial structures. Following the acquisition of the mask image, a contrast agent is injected in the carotid or vertebral artery providing enhanced visualization of the vascular anatomy. Once both CBCT acquisitions are completed, the imaging data is processed to generate a three-dimensional (3D) reconstruction of the intracranial vasculature after subtracting the first acquisition from the first.

The present study introduces a full pipeline to build a curvature-preserving planar representation of a such a 3D arborescence. A major challenge in vessel navigation algorithms is to avoid topological artifacts that are induced by the limitations of the voxel resolution. When arteries overlap in the image, large voxels may not allow us to distinguish between them accurately, causing “jumping arteries”. This occurs when the model incorrectly interprets overlapping or closely positioned arteries as a single vessel, causing the vessel representation to appear to jump or shift. This is the reason why we chose to base our model on high-resolution invasive imaging techniques, such as angiographies, rather than cross-sectional non-invasive images such as MRI or CT scans.

2.2 Pre-processing of the per operative 3D acquisition

Segmentation. As illustrated in Fig. 13.a, the raw data used in the present study consists of contrast-enhanced subtracted volumes of intra- and extra- cranial vessels. As an intermediate step to our 2-dimensional unfolding, we aim to convert this volume into a graph that represents the

geometry and topology of the cerebral vasculature. This is a non-trivial task, since a basic thresholding of the image (illustrated in Fig. 14.a) generates a binary mask or contour mesh with many topological defects that include fused vessels, disconnected components and noisy reconstruction artifacts.

To overcome these issues, we choose not to rely too much on the raw data intensity, which presents fluctuations along the vessels due to variations of the concentration of the contrast agent. Instead, we first apply a thresholding pass to extract a binary mask, and compute a truncated signed distance function to its boundary. This scalar field, illustrated in Fig. 13.b, takes values that range from 6mm (inside the thickest carotid artery) to -3mm (slightly outside of the vessel surface).

This operation is standard in image processing, but common implementations provided by e.g. the scikit-image library are too slow and memory hungry to scale up to large volumes of $500 \times 500 \times 500$ voxels on modest hardware. Instead, we choose to implement the Jump Flood Algorithm (Rong and Tan, 2006, 2007) using the Taichi library (Hu et al., 2019). This lets us focus on voxels that belong to the blood vessels without wasting computation time on the empty regions of the ambient space. We can then perform this distance computation in a couple of seconds on a typical laptop. More generally, our code relies extensively on the NumPy, SciPy and PyTorch libraries for computations (Harris et al., 2020; Virtanen et al., 2020; Paszke et al., 2017), as well as PyVista and ParaView for visualizations (Sullivan and Kaszynski, 2019; Ayachit, 2015).


From volume to curves. To extract the centerlines of the blood vessels, we then apply a Frangi vesselness filter (Frangi et al., 1998) on the signed distance volume. As in the previous step, we implement this operation using Taichi and recover the high-resolution vesselness volume of Fig. 13.c in a couple of seconds.

The scikit-image (van der Walt et al., 2014) and skan (Nunez-Iglesias et al., 2018) libraries then allow us to turn this volume into a set of connected 3D curves using an Otsu thresholding (Otsu et al., 1975), binary erosion, and a simple connectivity pattern between neighboring voxels. These steps rely on Lee’s method (Lee et al., 1994) to build the skeleton via a surface axis thinning algorithm that employs an octree data structure to examine the $3 \times 3 \times 3$ neighborhood surrounding each pixel. This iterative method progressively removes pixels while maintaining connectivity, until the binary skeleton volume converges to a stable state.


After applying a small Laplacian smoothing on the point coordinates to remove aliasing artifacts, we sample the distance volume of Fig. 13.b to estimate the pointwise vessel radii. Retaining only the largest component of the graph, we obtain the tubular model of Fig. 13.d.

2.3 Graphical representation of the vascular network


Enforcing a cycle-free topology. In order to produce an overlap-free 2D visualization of the brain artery network that preserves vessel lengths, we assume that our graph is a *tree*, without cycles. This is consistent with common knowledge about the brain anatomy: with the exception of the circle of Willis, a normal brain artery network should be cycle-free. Unfortunately, this strong topological assumptions does not hold in practice for the graph of Fig. 13. Due to limits in the voxel resolution, some vessels that run close to each other may be fused by our processing: our processing hallucinates “shortcuts” and topological cycles that do not correspond to realistic paths for endovascular interventions. Abnormal patient anatomies or vessel ruptures may also create cycles that should be taken into account, but that could be identified as explicit “outlier” edges on top of a normal tree structure.




(a) Raw image.




(b) Distance transform.




(c) Frangi vesselness filter.



(d) Graph with estimated vessel radius.



(e) Topological simplification.



(f) Geodesic distance to the root.

Figure 13: Main processing steps as we turn a raw 3D volume into a loop-free set of 3D curves.




Figure 14: Assessment of the quality of our graph reconstruction.

Pruning. Some topological denoising is thus required. To this end, and for the sake of computational efficiency, we first collapse all nodes with two neighbors in order to get a lighter graph that is illustrated in Fig. 13.e. This reduces the number of nodes from more than 25,000 to a mere 2,000 on average, while preserving the same topological structure. Then, we create a *root* node that represents the pumping heart and connect it to all nodes at the bottom end of the volume. Finally, we assign a cost to each *directed* edge of the graph with:

$$\text{weight}(a \rightarrow b) = \text{radius}(b) - \text{radius}(a) . \quad (3)$$

This weight encodes a sensible prior on artery trees, which is that the blood is likely to run from wide to narrow vessels. Using the Edmonds algorithm (Edmonds et al., 1967) provided by the `minimum_spanning_arborescence` method of the NetworkX library (Hagberg et al., 2008), we then find a cycle-free subgraph of the vascular network that connects all the nodes of the simplified network with minimal total weight, i.e. with minimal “upstream” flow going from thin to thick vessels. We display the result of this topological pruning on the high-resolution graph in Fig. 13.f. Cycle-inducing edges that have been pruned by the Edmonds algorithm are displayed in green, while the other ones are colored according to their geodesic distances to the root node.

Successes and limitations. As illustrated in Fig. 14.a-b, our method is able to capture most of the brain arterial network correctly. This includes challenging regions of the brain anatomy such as the two anterior cerebral arteries that run parallel to each other, between the left and right hemispheres, and are correctly modelled as distinct curves. In practice, we find that the topological structure of the main arteries (with radius larger than 2mm) is well preserved. On the other hand, two main sources of errors still remain:

- Some branches of the tree can be discarded, especially if they are thin and connected to the rest of the network via a sharp turn that is not highly contrasted in the raw 3D volume, as in the upper left corner of Fig. 14.b.
- Although our method deals effectively with vessels that run parallel to each other, it does not handle “crossings” well. This is evident in Fig. 14.c, and is a consequence of the fact that our topological cleaning procedure can only delete graph edges and is unable to split a “crossroad” node with 4 neighbors into two separate curves.

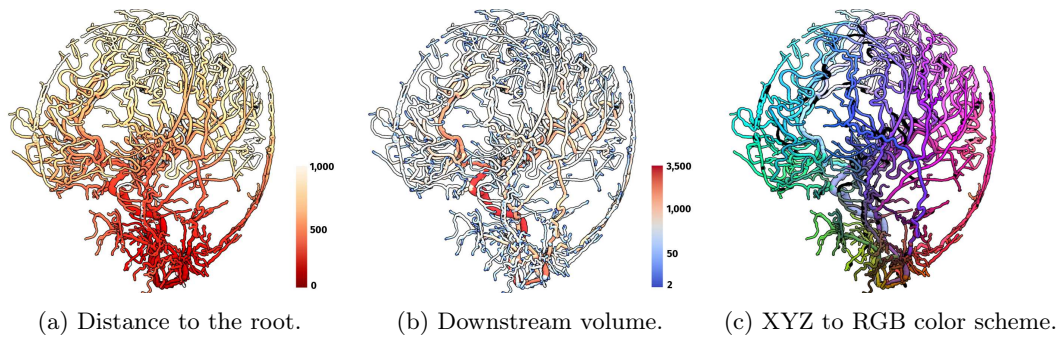


Figure 15: Point features on our graph representation.





Figure 16: From 3D to 2D representations.

Fortunately, these segmentation artifacts only affect a tiny portion of the brain artery network. Errors typically occur in regions that are too remote to be accessible with current endovascular devices anyway. Nevertheless, getting a reliable and precise network topology from our 3D images is one of our main targets for future works.

Point features. For the sake of visualization, we enrich our graph with node features that are computed recursively:

- The geodesic distance to the root, which is illustrated in Fig. 13.f and Fig. 15.a.
- The volume of downstream vessels, which is is illustrated in Fig. 15.b and computed by identifying each edge with a truncated cone.
- The XYZ coordinates of the barycenter of downstream vessels, that we display directly as RGB colors in Fig. 15.c and Fig. 16.

2.4 Planar graphical representation of the vascular network

Planar embedding of a graph. Over the past decades, many algorithms have been proposed to draw graphs on a 2D plane (Battista et al., 1994). As illustrated in Fig. 17, standard methods such as the Kamada–Kawai algorithm (Kamada and Kawai, 1989) force us to choose between preserving path lengths or providing intersection-free visualizations. Although significant progress has recently been made in the untangling of 3D knots and curves (Yu et al., 2021), we found that existing force-directed drawing methods could not reliably display in 2D the complex topologies of brain arterial networks.

In order to satisfy all of our constraints, we thus propose an original method for planar graph representation that is tailored for the requirements of interventional radiologists. We aim at producing aesthetically pleasing drawings of the neurovasculature that improve the graph readability, without compromising anatomical considerations.

Recursive tree embedding. As a first step, we compute a planar representation of the tree using a recursive algorithm:

1. The root node a is placed at the origin of the 2D plane $(x(a), y(a)) = (0, 0)$ and is allocated an angular “budget” or interval $[\theta_{\min}(a), \theta_{\max}(a)]$ that is equal to a full half-plane $[0, \pi)$.
2. Consider the K children b_1, \dots, b_K of node a , and sort them according to the position of the downstream barycenter along the front-to-back axis.
3. Compute the relative weight $w_k \in [0, 1]$ of each child b_k as the vessel volume downstream of edge $a \rightarrow b_k$ divided by the total volume downstream of node a . By construction, the weights w_k sum up to 1 and correspond to the relative importance of each branch flowing out of a .
4. Given an angle budget $[\theta_{\min}(a), \theta_{\max}(a)]$, split it into K consecutive intervals:

$$[\theta_{\min}(b_1), \theta_{\max}(b_1)], \dots, [\theta_{\min}(b_K), \theta_{\max}(b_K)] \quad (4)$$

with lengths that are proportional to the weights w_1, \dots, w_K .




Figure 17: Kamada-Kawai and log-linear representations of the vascular tree.

5. Define the angle $\theta_{a \rightarrow b_k}$ as the midpoint $\frac{1}{2}(\theta_{\min}(b_k) + \theta_{\max}(b_k))$, and the length $\ell(a \rightarrow b_k)$ as the Euclidean length of the (small) corresponding segment in the original 3D graph.
6. Define the 2D positions of the children nodes with:

$$\begin{pmatrix} x(b_k) \\ y(b_k) \end{pmatrix} = \begin{pmatrix} x(a) + \ell(a \rightarrow b_k) \cos(\theta(a \rightarrow b_k)) \\ y(a) + \ell(a \rightarrow b_k) \sin(\theta(a \rightarrow b_k)) \end{pmatrix}. \quad (5)$$

7. Run the algorithm recursively on the children b_k , with corresponding positions $(x(b_k), y(b_k))$ and angular budgets $[\theta_{\min}(b_k), \theta_{\max}(b_k)]$.

We perform this embedding at a high resolution, with the understanding that most nodes have only one child and thus correspond to straight junctions between small articulated segments. By construction, this method preserves the geodesic lengths of the blood vessels and creates no overlap between branches. It also lets us sort the branches according to an arbitrary criterion – in our case, along the front-to-back axis as in a slice along the sagittal plane. On the other hand, as illustrated in the left panel of Fig. 16.b, it hardly allows a radiologist to identify individual vessels at a glance.

Electrostatic repulsion. To improve the readability of our representation, we then relax the constraints on the angles $\theta(a \rightarrow b)$ and progressively move the branches away from each other. Using a modified electrostatic kernel:

$$k(a, b) = \frac{10}{\|a - b\|_{\mathbb{R}^2} + \varepsilon} - \|a - b\| \quad (6)$$

with gradient:

$$k'(a, b) = (a - b) \cdot \left(\frac{10}{(\|a - b\|_{\mathbb{R}^2} + \varepsilon)^3} + \frac{1}{\|a - b\|} \right), \quad (7)$$

a simple electrostatic repulsion model is to consider, on each node a , a force:

$$\text{Force}(a) = - \sum_{\text{nodes } b} \text{charge}(b) k'(a, b) \in \mathbb{R}^2. \quad (8)$$

For the sake of performance, we restrict the sum to the 100-nearest neighbors b of each node a and use constant charges equal to 1/100. As highlighted in (Yu et al., 2021), optimizing such an energy with respect to node positions is a surprisingly complex endeavour. Instead, we choose to act on the edge angles $\theta_{a \rightarrow b}$ with updates:

$$\theta_{a \rightarrow b} \leftarrow \theta_{a \rightarrow b} + \frac{\Delta t}{\ell(a \rightarrow b)} \begin{pmatrix} -\sin(\theta_{a \rightarrow b}) \\ +\cos(\theta_{a \rightarrow b}) \end{pmatrix} \cdot \frac{1}{\#C(b)} \sum_{c \in C(b)} \text{Force}(c), \quad (9)$$

where $C(b)$ denotes the set of all nodes downstream from b , including b itself, and “ \cdot ” denotes a dot product between a 2D vector that is orthogonal to the edge $(a \rightarrow b)$ and the average force downstream from b . This optimization on the edge angles guarantees that edge lengths are preserved at all steps of the algorithm. Using a small enough time step $\Delta t = 0.1$ to avoid collisions between the 2D vessels, we then obtain at convergence the embedding that is displayed in the middle of Fig. 16.b.

Displaying the curvature. As a final requirement, we bend the vessels to best reflect their local curvature. To this end, we first fit a plane to each point of the 3D arterial network by performing a Principal Component Analysis on geodesic neighborhoods of radius $r = 100\text{mm}$. Starting from the root, we orient these planes consistently and compute, for each pair of consecutive 3D edges $(a \rightarrow b)$ and $(b \rightarrow c)$, an oriented bending angle $\phi_{a \rightarrow b \rightarrow c}$ in the local 2D plane at point b .

This bending corresponds to a target curvature for our 2D embeddings. Starting from an angle 0 at the root, we integrate it along the tree edges to obtain a target orientation $\hat{\theta}_{a \rightarrow b}$ for each 2D edge. Then, denoting by “Smooth” a Gaussian smoothing operator of radius $r = 100\text{mm}$ on the graph, we compute a new bending force for each edge $(a \rightarrow b)$ as:

$$\text{Bending}(a \rightarrow b) = \theta_{a \rightarrow b} - \hat{\theta}_{a \rightarrow b} - \text{Smooth}(\theta - \hat{\theta})_{a \rightarrow b}. \quad (10)$$

The smoothing operator is required to avoid too much drift in the bending formula. To balance this new bending force with the electrostatic repulsion, we then simply iterate until convergence the following update on the angles of the 2D embedding edges:

$$\theta_{a \rightarrow b} \leftarrow \theta_{a \rightarrow b} + \frac{\Delta t}{\ell(a \rightarrow b)} \begin{pmatrix} -\sin(\theta_{a \rightarrow b}) \\ +\cos(\theta_{a \rightarrow b}) \end{pmatrix} \cdot \frac{1}{\#C(b)} \sum_{c \in C(b)} \text{Force}(c) - \kappa \text{Bending}(a \rightarrow b), \quad (11)$$

where $\kappa = 0.002$ is a small scaling factor. We end up with the planar representation in the right panel of Fig. 16.b, that we find satisfying for our purposes. Going forward, we intend to work on color codes that will highlight accessible zones for different types of endovascular devices, and extend our method to other organs.

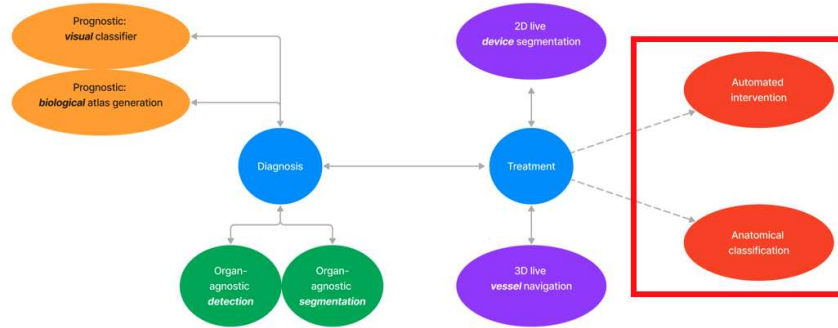


Figure 18: Red projects

3 Extension towards the origins of the vascular network: live embryology

This section refers to the *red projects*. Special thanks to Pr Alain Chédotal, Eimad Shotar and Armelle Lesaunier for this ongoing project.

3.1 Introduction

The expansion of endovascular therapies is driven by technological advancements in imaging and navigation tools. Additionally, a **deeper understanding of the anatomy relevant to vascular procedures** has contributed significantly to these advancements, enabling procedures to be tailored to individual anatomies and reducing the risk of complications.

While new therapeutic procedures shed light on descriptive anatomy, comparative and evolutionary anatomy offers unique insights into human anatomy. **Embryology** also contributes to our **understanding of adult anatomy**. In the neurovascular field, the anatomical complexity of the anterior communicating artery complex is a classic example of the interplay between embryology, anatomy, and pathology. Developmental anomalies in this complex are linked to genetic polymorphisms identified in animal models, major risk factors for developing anterior communicating artery aneurysms, and significant variations in tolerance to pathological or therapeutic internal carotid artery occlusions.

Recent years have seen significant progress in understanding the **microvascular embryology** of the central nervous system, particularly regarding the regulation of angiogenesis from perineural to intraneural vascular plexuses. This is partly due to the emerging concept of the "neurovascular link," suggesting shared guidance mechanisms between angiogenesis and axonal growth.

While embryology has shed light on human anatomy, the macroscopic neurovascular anatomy remains relatively underexplored. Most of our knowledge in this area comes from Dorcas Hager Padget's articles published between 1948 and 1957, which have had a significant influence. However, there are limitations to this knowledge, as it was purely descriptive and lacked information on morphogenesis control. Additionally, the methodology of embryonic sectioning or dissection may introduce deformations and biases in three-dimensional reconstructions.

The advent of endovascular therapies, the need for anatomical knowledge, and the development

of new embryological tools provide an opportunity to reevaluate and enrich our understanding of neurovascular morphogenesis.

Notably, recent advances in imaging techniques allow **the exploration of entire organs at cellular resolution during organogenesis, without the need to destroy the tissue**. These techniques involve rendering biological samples **transparent through various methods like delipidation, decolorization, and decalcification**, which adjust the refractive index of the sample to match its surroundings. These methods, combined with effective immunolabeling and volumetric microscopy, enable the examination of entire organs, even whole rodent bodies.

3.2 Dataset and study objectives

Methods like **3DISCO (Three-dimensional imaging of solvent cleared organs) and its derivatives iDISCO (immunolabeling-enabled three-dimensional imaging of solvent-cleared organs), iDISCO+, uDISCO, or vDISCO are compatible with fluorescent immunolabeling of intact/whole organs and embryos**. They make tissue optically transparent, allowing the visualization of the microscopic organization of tissues throughout their thickness using a high-resolution laser scanning microscope. This approach presents a unique opportunity to explore morphogenesis and has already been applied to various animal models to study axonal guidance in mouse embryos, gonadal, renal, and retinal morphogenesis, and the development of the enteric nervous system in chicken embryos. It has also provided an unparalleled description of the development of many human embryonic organs.

The method was detailed in Cell Blain et al. (2023), detailing early developments of human head embryos from over seventy specimens. The method included staining using conjugated antibodies for labeling large human embryo and fetal samples and iDISCO/ iDISCO+ clearing. This study detailed the development of the head skeleton, head muscles and innervation, salivary and lacrimal glands, and head and neck arteries. No automatic classification process was proposed.

We are particularly interested in the **arterial** segmentation of different embryos. The staining process used immunostaining for smooth muscle actin (SMA). The authors note that "Last, to facilitate the visualization of fully annotated 3D image datasets and their use for outreach and teaching purposes, inter- active 3D models of embryonic head arteries and skeleton were created by exporting segmented meshes into Verge3D, a web- based interface and toolkit".

We aim to reproduce these results for the pelvic arteries using the planar representation of the 3D vasculature described in the previous section. Using vessel-specific coating, embryo vasculature was manually segmented and saved as blocks in a multiblock dataset. Below is portrayed the specific segmentation task on pelvic arteries, an ongoing multidisciplinary project with the Institut de la Vision (France).

In this ongoing collaboration with the Institut de la Vision, we aim to apply our solution for generating planar representations of newly portrayed ultra high resolution vascular segmentation at early stages of human development. We hope to generate new 'classifications' with higher legibility while preserving important morphological parameters, as explained earlier.




Figure 19: From Blain et al, analysis of developing cephalic arteries in human embryos, *immunostained for SMA*




Figure 20: **3D analysis of developing rectal arteries in human embryos, courtesy of Armelle Lesaunier.** All pictures are light sheet florescence microscopy images of solvent-cleared immunostained human embryos at various stages of development. Embryos were immunostained with SMA or CD31 antibodies. (A-D) Isolated segmented pelvic arterial vasculature. (A) shows a lateral view of a PCW5.6 embryo. (B) shows a lateral view of a PCW7 embryo. (C) shows a lateral view of a PCW 8.5 embryo. (D) shows a frontal view of a PCW11 embryo. The panel illustrate the evolution of the rectal arterial blood supply from two main supplies (superior rectal artery and sacral rectal arteries) to the adult pattern with persistent sacral rectal arteries. Ao, Aorta; IMA, Inferior mesenteric artery; ISA, Interssegmental arteries; MSA, Median sacral artery; ScRA, Sacral rectal arteries; MRA, Middle rectal arteries; IRA, Inferior rectal arteries; OA, Ombilical arteries; PA, Pudendal arteries.

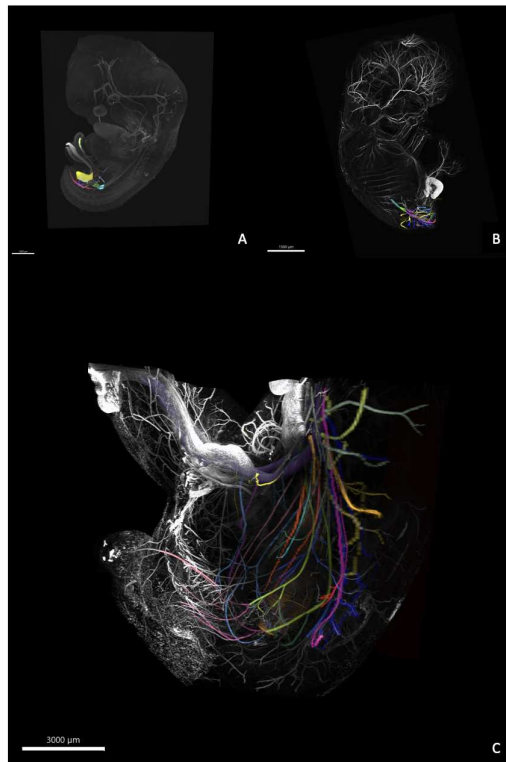


Figure 21: **Courtesy of Armelle Lesaunier.** Light-sheet fluorescent microscopy images of solvent-cleared immunostained human embryos at various stages of development. All panels show segmented pelvic arteries and the raw signal of the immunostained embryos. (A) is a lateral view of a PCW5.6 embryo immunostained with SMA and CD31 antibodies. (B) is a lateral view of a PCW8.5 embryo immunostained with SMA and Col2. (C) is a frontal oblique left view of a PCW11 embryo immunostained with SMA.

References

- Utkarsh Ayachit. *The paraview guide: a parallel visualization application*. Kitware, Inc., 2015.
- Maxime Barat, Anthony Dohan, Maureen Kohi, Clement Marcelin, Jean-Pierre Pelage, Alban Denys, Sebastian Mafeld, Claire S Kaufman, Philippe Soyer, and Francois H Cornelis. Treatment of adenomyosis, abdominal wall endometriosis and uterine leiomyoma with interventional radiology: A review of current evidences. *Diagn. Interv. Imaging*, 105(3):87–96, March 2024.
- Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G Tollis. Algorithms for drawing graphs: an annotated bibliography. *Comput. Geom.*, 4(5):235–282, October 1994.
- Raphael Blain, Gérard Couly, Eimad Shotar, Joséphine Blévin, Maryne Toupin, Anais Favre, Ali Abjagh, Megumi Inoue, Edwin Hernández-Garzón, Frédéric Clarençon, Frédéric Chalmel, Séverine Mazaud-Guittot, Paolo Giacobini, Yorick Gitton, and Alain Chédotal. A tridimensional atlas of the developing human head. *Cell*, 186(26):5910–5924.e17, December 2023.
- Tom Boeken, Carole Dean, Olivier Pellerin, and Marc Sapoval. How artificial intelligence will reshape our interventional units. *CardioVascular and Interventional Radiology*, 47(2):283–284, October 2023a. ISSN 1432-086X. doi: 10.1007/s00270-023-03590-z. URL <http://dx.doi.org/10.1007/s00270-023-03590-z>.
- Tom Boeken, Jean Feydy, Augustin Lecler, Philippe Soyer, Antoine Feydy, Maxime Barat, and Loïc Duron. Artificial intelligence in diagnostic and interventional radiology: Where are we now? *Diagn. Interv. Imaging*, 104(1):1–5, January 2023b.
- Xu Chen, Xiangde Luo, Yitian Zhao, Shaoting Zhang, Guotai Wang, and Yalin Zheng. Learning euler’s elastica model for medical image segmentation, 2020. URL <https://arxiv.org/abs/2011.00526>.
- Taco S. Cohen and Max Welling. Group equivariant convolutional networks. 2016. doi: 10.48550/ARXIV.1602.07576. URL <https://arxiv.org/abs/1602.07576>.
- William Crinnion, Ben Jackson, Avnish Sood, Jeremy Lynch, Christos Bergeles, Hongbin Liu, Kawal Rhode, Vitor Mendes Pereira, and Thomas C Booth. Robotics in neurointerventional surgery: a systematic review of the literature. *J. Neurointerv. Surg.*, pages neurintsurg–2021–018096, November 2021.
- Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, July 1945.
- Jack Edmonds et al. Optimum branchings. *Journal of Research of the national Bureau of Standards B*, 71(4):233–240, 1967.
- Pepe Eulzer, Monique Meuschke, Gabriel Mistelbauer, and Kai Lawonn. Vessel maps: A survey of map-like visualizations of the cardiovascular system. In *Computer Graphics Forum*, volume 41, pages 645–673. Wiley Online Library, 2022.
- Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI’98: First International Conference Cambridge, MA, USA, October 11–13, 1998 Proceedings 1*, pages 130–137. Springer, 1998.

- Rahul Ghosh, Kelvin Wong, Yi Jonathan Zhang, Gavin W Britz, and Stephen T C Wong. Automated catheter segmentation and tip detection in cerebral angiography with topology-aware geometric deep learning. *J. Neurointerv. Surg.*, pages jnis-2023-020300, June 2023.
- Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.
- Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédo Durand. Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics (TOG)*, 38(6):1–16, 2019.
- Gokmen Kahilogullari, Hasan Caglar Ugur, Ayhan Comert, Ibrahim Tekdemir, and Yucel Kanpolat. The branching pattern of the middle cerebral artery: is the intermediate trunk real or not? an anatomical study correlating with simple angiography. *Journal of neurosurgery*, 116(5):1024–1034, 2012.
- Tomihisa Kamada and Satoru Kawai. An algorithm for drawing general undirected graphs. *Inf. Process. Lett.*, 31(1):7–15, April 1989.
- T C Lee, R L Kashyap, and C N Chu. Building skeleton models via 3-D medial surface axis thinning algorithms. *CVGIP Graph. Models Image Process.*, 56(6):462–478, November 1994.
- Arash Najafi, Roberto Luigi Cazzato, Bernhard C Meyer, Philippe L Pereira, Angel Alberich, Antonio López, Maxime Ronot, Jan Fritz, Monique Maas, Sean Benson, Patrick Haage, and Fernando Gomez Munoz. CIRSE position paper on artificial intelligence in interventional radiology. *Cardiovasc. Radiol.*, 46(10):1303–1307, October 2023.
- Juan Nunez-Iglesias, Adam J Blanch, Oliver Looker, Matthew W Dixon, and Leann Tilley. A new python library to analyse skeleton images confirms malaria parasite remodelling of the red blood cell membrane skeleton. *PeerJ*, 6:e4312, 2018.
- Nobuyuki Otsu et al. A threshold selection method from gray-level histograms. *Automatica*, 11 (285-296):23–27, 1975.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Guodong Rong and Tiow-Seng Tan. Jump flooding in GPU with applications to Voronoi diagram and distance transform. In *Proceedings of the 2006 symposium on Interactive 3D graphics and games*, pages 109–116, 2006.
- Guodong Rong and Tiow-Seng Tan. Variants of jump flooding algorithm for computing discrete Voronoi diagrams. In *4th international symposium on voronoi diagrams in science and engineering (ISVD 2007)*, pages 176–181. IEEE, 2007.

- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. 2015.
- Yuya Sakakura, Kenichi Kono, and Takeshi Fujimoto. Real time artificial intelligence assisted carotid artery stenting: a preliminary experience. *Journal of NeuroInterventional Surgery*, pages jnis-2024-021600, April 2024. ISSN 1759-8486. doi: 10.1136/jnis-2024-021600. URL <http://dx.doi.org/10.1136/jnis-2024-021600>.
- Suprosanna Shit, Johannes C Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylyka, Josien P W Pluim, Ulrich Bauer, and Bjoern H Menze. CIDice - a novel topology-preserving loss function for tubular structure segmentation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2021.
- T. Sørensen. *A Method of Establishing Groups of Equal Amplitude in Plant Sociology Based on Similarity of Species Content and Its Application to Analyses of the Vegetation on Danish Commons*. Biologiske skrifter. Munksgaard in Komm., 1948. URL <https://books.google.fr/books?id=rpS8GAAACAAJ>.
- Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, Lecture notes in computer science, pages 240–248. Springer International Publishing, Cham, 2017.
- C Sullivan and Alexander Kaszynski. PyVista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (VTK). *Journal of Open Source Software*, 4(37):1450, 2019.
- Stéfan van der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, Tony Yu, and the scikit-image contributors. scikit-image: image processing in Python. *PeerJ*, 2:e453, 6 2014. ISSN 2167-8359. doi: 10.7717/peerj.453. URL <https://doi.org/10.7717/peerj.453>.
- Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. SciPy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.
- Chris Yu, Henrik Schumacher, and Keenan Crane. Repulsive curves. *ACM Trans. Graph.*, 40(2), 2021.
- Guifang Zhang, Hon-Cheng Wong, Cheng Wang, Jianjun Zhu, Ligong Lu, and Gaojun Teng. A temporary transformer network for guide-wire segmentation. In *2021 14th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*. IEEE, October 2021.
- Guifang Zhang, Hon-Cheng Wong, Jianjun Zhu, Tao An, and Cheng Wang. Jigsaw training-based background reverse attention transformer network for guidewire segmentation. *Int. J. Comput. Assist. Radiol. Surg.*, 18(4):653–661, December 2022.

Supplementary Material: French Summary

August 15, 2024

Contents

1	Partie I. Introduction	2
2	Partie II. Détection d'anomalies focales en utilisant le LDDMM	3
3	Partie III. Utilisation d'un marqueur biologique pour prédire la réponse tumorale	5
4	Partie IV. Représentation planaire de l'arborescence vasculaire 3D	8

1 Partie I. Introduction

L'intégration de la vision dans les interventions guidées par l'image a le potentiel de transformer notre pratique médicale. Ce travail pose les bases pour l'avenir des interventions autonomes dans notre domaine spécifique concernant les patients atteints de cancer, en abordant des composants clés nécessaires à sa réalisation. Nous explorons tout d'abord l'impact transformateur de l'intelligence artificielle sur les capacités physiques des radiologues interventionnels. Nous soulignons la nécessité de relever les défis techniques et éthiques. La collaboration interdisciplinaire et des processus d'évaluation rigoureux sont mis en avant comme essentiels pour une intégration sûre de l'IA dans la pratique clinique.

Nous proposons une méthode pour détecter les anomalies focales sur les images volumétriques en coupe. En tirant parti du cadre du Large Diffeomorphic Deformation Metric Mapping (LDDMM), cette approche démontre une reconstruction améliorée des objets et une localisation précise des lésions. Dans le même cadre, nous proposons un classificateur, où la sélection des patients présente des défis uniques en raison des rapports bénéfice/risque complexes. Pour aller au-delà des images, les données cliniques issues de l'analyse de l'ADN tumoral sont intégrées dans une étude prospective spécifiquement menée pour ce travail.

Les réseaux antagonistes génératifs (GAN) et la modélisation des atlas en utilisant les algorithmes Markov Chain Monte Carlo - Stochastic Approximation Expectation-Maximization (MCMC-SAEM) sont utilisés pour prédire les trajectoires des patients. Cette approche permet d'explorer de nouvelles trajectoires, améliorant ainsi notre compréhension de la progression de la maladie et de la réponse au traitement en relation avec l'ADN tumoral circulant.

Enfin, nous explorons des techniques avancées de visualisation pour la vascularisation 3D in vivo et ex vivo. Nous proposons une représentation plane d'une anatomie 3D, offrant une voie prometteuse pour une exploration et une compréhension plus approfondies. Ensemble, ces sections offrent des solutions pour la réalisation des interventions autonomes dans notre domaine.

Les raisons qui précèdent ce travail sont multiples. Aujourd'hui, l'intégration de l'IA et de la robotique pourrait encore propulser la médecine interventionnelle vers de nouvelles applications et de meilleurs résultats pour les patients. Bien que les bénéfices cliniques des procédures assistées par robot soient multiples, y compris une meilleure précision et des temps de procédure réduits, des limitations persistent, comme l'absence de retour tactile complet. La pleine autonomie des robots chirurgicaux nécessitera des avancées majeures en IA, en vision par ordinateur et en technologies de capteurs.

Ce travail se concentre sur les prérequis pour une automatisation complète, notamment la sélection des patients avant l'intervention et leur traitement avec les dispositifs actuels. Il explore deux approches principales en radiologie interventionnelle : les approches endovasculaires et percutanées, tout en abordant les défis éthiques et techniques associés à l'intégration de l'IA. La collaboration interdisciplinaire, les tests précliniques rigoureux, et la formation adaptée seront essentiels pour une intégration sûre et efficace de ces systèmes dans la pratique clinique.

Le but de ce travail est de montrer comment les avancées récentes en mathématiques appliquées permettent de poser les bases pour une radiologie interventionnelle autonome, avec un focus particulier sur les patients atteints de cancer colorectal, depuis leur sélection jusqu'aux thérapies guidées en 2D et 3D en direct. Les projets verts se concentrent sur la détection et la segmentation des tumeurs à partir des scanners pré-thérapeutiques, les projets orange sur les biomarqueurs prédictifs, et les projets violets sur la compréhension des images en temps réel pour l'autonomie future.

2 Partie II. Détection d’anomalies focales en utilisant le LDDMM

Cette section présente une méthode pour détecter les anomalies focales sur des images volumétriques en coupe transversale sans nécessiter d’annotations ou de jeux de données d’entraînement étendus.

En s’appuyant sur le cadre du Large Diffeomorphic Deformation Metric Mapping (LDDMM), cette méthode est illustrée à l’aide de scanners du foie et d’IRM cérébrales. La présence d’anomalies focales nous empêche d’appliquer directement une déformation d’un modèle vers un patient avec un nombre inconnu d’anomalies focales. Nous proposons d’utiliser cette contrainte diffeomorphe à notre avantage en modélisant les résidus (c’est-à-dire la différence entre le modèle déformé et l’observation) comme une matrice clairsemée (sparse matrix), en plus d’un bruit indépendant.

Nous pouvons utiliser les résidus de la déformation diffeomorphe à partir d’un modèle de contrôle pour détecter et segmenter les lésions dans n’importe quel organe. Ce qui ne peut être reconstruit comme une déformation diffeomorphe du modèle est ainsi placé dans cette matrice et classé comme une anomalie.

Nous démontrons que l’estimation *simultanée* des déformations et des anomalies améliore les résultats en minimisant les erreurs de reconstruction. Appliquée à des données simulées, notre méthode permet une meilleure reconstruction des objets et une localisation précise des lésions. Cette adaptabilité s’étend aux données réelles, où nous appliquons l’approche aux ensembles de données du foie et du cerveau. Dans le jeu de données du foie, notre approche a détecté 94 % des métastases, démontrant ainsi son efficacité.

Plus précisément, cette étude présente une approche géométrique pour détecter et segmenter des anomalies focales dans les organes à partir d’images médicales volumétriques. Contrairement aux méthodes traditionnelles, cette approche ne nécessite ni annotations ni grands ensembles de données d’entraînement.

Classiquement, les méthodes de segmentation basées sur l’apprentissage profond, telles que les U-Nets, ont contribué aux besoins spécifiques des données biomédicales, mais nécessitent souvent de grands ensembles de données annotées. Cependant, les solutions généralisables, capables de bien fonctionner sur plusieurs tâches sans nécessiter de modèles spécifiques, gagnent en importance. Cette étude explore donc la possibilité de segmenter les anomalies focales à partir de petits ensembles de données, en utilisant un cadre géométrique dans un modèle de détection agnostique aux organes.

En pratique, la méthode proposée utilise un modèle de contrôle basé sur des patients sans anomalies pour créer un modèle de référence. Les anomalies sont alors détectées comme des éléments qui ne peuvent être reconstruits par une déformation diffeomorphe du modèle de contrôle. Cette approche a été testée sur des ensembles de données simulées et réelles, montrant une meilleure reconstruction des objets et une localisation précise des lésions. Des applications ont été réalisées sur un ensemble de données de métastases hépatiques et un ensemble de données d’IRM cérébrales pour illustrer la polyvalence et l’efficacité de cette méthode.

Dans cette section, nous avons démontré l’efficacité de l’utilisation des résidus dérivés de la déformation diffeomorphe d’un modèle de contrôle pour détecter et segmenter les lésions au sein d’un organe donné. Nos résultats indiquent que cette approche facilite non seulement l’identification des lésions, mais améliore également la reconstruction diffeomorphe globale des données observées. Un avantage inhérent à notre méthode réside dans son indépendance vis-à-vis de vastes ensembles de données patient ou d’annotations de professionnels de la santé, ce qui la rend particulièrement adaptée à des situations où l’acquisition de grandes bases de données est difficile, comme dans certains protocoles de traitement spécifiques.

Nos expérimentations ont consisté à évaluer la performance de cette méthodologie sur des ensembles de données comprenant des cerveaux atteints de gliome et des foies présentant des métastases du cancer colorectal. Notamment, dans le cas des cerveaux atteints de gliome, notre algorithme a réussi à enregistrer les gyri et les ventricules tout en identifiant avec précision les tumeurs dans la matrice d'anomalies.

Cependant, plusieurs défis subsistent et méritent une attention particulière. Premièrement, bien que notre méthode détecte efficacement les tumeurs, elle capture également de petites erreurs de reconstruction. Pour y remédier, des étapes de post-traitement sont nécessaires pour isoler et extraire uniquement les lésions pertinentes. Une première option consiste à utiliser des opérations de fermeture, bien que cette approche comporte le risque potentiel d'éliminer par inadvertance de plus petites anomalies lors du post-traitement. De plus, une stratégie plus affinée pour gérer les erreurs de reconstruction est impérative, car l'absence d'une tumeur dans l'objet reconstruit peut entraîner son exclusion accidentelle. Un affinage supplémentaire via le post-traitement de la matrice d'anomalies est essentiel pour récupérer les anomalies éventuellement négligées et garantir la précision de notre méthode de détection.

3 Partie III. Utilisation d'un marqueur biologique pour prédire la réponse tumorale

Un nouveau biomarqueur biologique a émergé comme une percée dans la gestion du cancer colorectal : l'ADN tumoral circulant (ADNc). L'ADNc est constitué d'ADN provenant de cellules tumorales circulant dans le sang. Détecté en quantités infimes grâce à des techniques de séquençage, ce biomarqueur permet une estimation plus sensible et précoce de la charge tumorale par rapport aux méthodes d'imagerie conventionnelles. Sa pertinence pronostique a été démontrée non seulement dans le cancer colorectal non métastatique, mais aussi dans la sélection des patients pour une chirurgie dans les cas de maladie oligométastatique.

Notre projet prospectif complémentaire à cette thèse visait à clarifier la valeur pronostique et thérapeutique de l'analyse répétée de l'ADNc avant et après les traitements par des radiologies interventionnelles, tels que l'ablation. Nous avons décidé de collecter et d'analyser des échantillons de nos patients, et nous pensons être les premiers à utiliser l'ADNc dans un tel contexte.

La radiologie interventionnelle offre une alternative non chirurgicale par traitement percutané. Dans ce projet, seule l'ablation thermique (utilisant des micro-ondes ou des radiofréquences pour la destruction par la chaleur) est envisagée. Les données sur l'ablation thermique des métastases hépatiques, issues d'essais randomisés majeurs, ont positionné cette technique dans l'arsenal thérapeutique. Ils ont montré une survie globale accrue chez les patients traités par ablation thermique par rapport au traitement systémique seul. Ce traitement a retardé la progression et guéri une fraction des patients, bien que la moitié des patients traités aient finalement succombé à leur maladie dans les cinq ans suivant le traitement. En comparaison, après une chirurgie hépatique, 70 % des patients traités connaissent une récurrence ou une progression dans les deux ans. Cependant, la place optimale de l'ablation thermique à visée curative dans le parcours oncologique du patient reste à définir, notamment en relation avec le contrôle global de la maladie par les traitements systémiques.

Les équipes d'oncologie et de biologie de l'HEGP sont des expertes dans le domaine de l'ADNc depuis de nombreuses années, notamment à travers le projet Carpem. Cette technique recherche de manière exhaustive les anomalies moléculaires liées aux gènes étudiés. Les panels ciblent des gènes spécifiques à la maladie ; leur coût raisonnable et les résultats rapides permettent une utilisation clinique de routine à l'HEGP. En effet, l'ADNc est mesuré de manière routinière à l'HEGP dans les cas de progression ou de stabilité de la maladie ambiguë.

Un tel biomarqueur serait particulièrement adapté à la situation clinique des patients subissant une ablation thermique hépatique. De nouveaux biomarqueurs - ADN tumoral circulant, cellules tumorales circulantes, vésicules extracellulaires circulantes, etc. - ont introduit le concept de maladie résiduelle minimale, et plus largement, de nouvelles définitions de la maladie oligométastatique. Au-delà de la sémantique, ces nouvelles entités cliniques ont des implications thérapeutiques. En fonction du stade (d'une seule métastase à une maladie multi-métastatique), les options de traitement évoluent parallèlement aux avancées des techniques ablatives locales dans le foie (ablation thermique, électroporation, radiothérapie, thérapies intra-artérielles, et chirurgie mini-invasive).

L'ADN tumoral circulant (ADNc) a été étudié dans plusieurs contextes tout au long de l'histoire oncologique du cancer colorectal. Dans le cancer colorectal non métastatique, les publications s'accordent sur la valeur pronostique de ce biomarqueur (Loft et al., 2023) : les patients avec de l'ADNc positif avant ou après la chirurgie présentent un pronostic moins favorable que les patients sans ADN tumoral circulant. La présence d'ADNc après la chirurgie reflète une maladie résiduelle et place le patient dans un groupe à haut risque de récurrence précoce. De même, dans la maladie

métastatique, la persistance de l'ADNc reflète une maladie insuffisamment traitée. Il est détecté dans plus de 90 % des cas.

Naturellement, les équipes chirurgicales ont examiné l'ADNc dans le cadre périopératoire ou dans la maladie potentiellement résecable. Les résultats de ces séries chirurgicales ne peuvent pas être extrapolés à l'ablation thermique en raison de mécanismes biologiques et moléculaires différents, même si l'objectif final reste l'élimination tumorale (c'est-à-dire la réduction complète de la charge tumorale et la détection négative de l'ADNc). Lors de l'ablation thermique, la tumeur reste en place, et l'énergie thermique induit la nécrose, l'ischémie, l'apoptose et/ou la mort cellulaire immunogène, générant une immunogénicité intrinsèque (concept de "priming").

Les données publiées sur l'ablation thermique et l'ADNc sont très limitées. Une étude préliminaire incluant des patients traités par ablation thermique et radiothérapie a montré que tous les cinq patients avec de l'ADNc positif après l'opération avaient un pronostic défavorable. À ce jour, aucune publication n'a abordé les différents profils évolutifs de l'ADNc après une ablation thermique. C'est pourquoi nous avons décidé d'ajouter cette caractéristique à notre travail sur la sélection des patients.

Il s'agit d'une étude observationnelle prospective monocentrique à un seul bras, approuvée par le comité d'éthique IRB #00011928 (CERAPHP Centre 2023-01-05). Tous les patients consécutifs traités entre novembre 2021 et juin 2022 pour des métastases hépatiques, pulmonaires ou osseuses du cancer colorectal, référés par la commission multidisciplinaire des tumeurs pour une ablation thermique à visée curative, ont été sélectionnés.

Les critères d'inclusion comprenaient les métastases du cancer colorectal accessibles à une ablation thermique à visée curative, telles qu'évaluées par la commission des tumeurs, une charge tumorale mesurable (méthode RECIST 1.1), un âge supérieur à 18 ans, un score ECOG de 0 ou 1, une maladie stable sous traitement systémique, un suivi réalisable, et une tumeur colorectale primitive réséquée. Les critères d'exclusion étaient un traitement combiné avec une chirurgie ou une radiothérapie, une tumeur primitive non réséquée, une fonction hépatique, pulmonaire ou rénale altérée, un refus de participer à l'étude, et la présence concomitante d'un autre cancer non colorectal.

Les interventions étaient planifiées sous anesthésie générale lors d'une hospitalisation de deux jours, et le bevacizumab était interrompu pendant au moins trois semaines avant toute procédure interventionnelle. La taille, la localisation, le profil moléculaire de la tumeur, la localisation de la tumeur primitive, les types et le nombre de lignes de traitement antérieures, les marges d'ablation et les traitements post-interventionnels ont été enregistrés. L'ablation thermique comprenait l'ablation par micro-ondes pour les lésions hépatiques ou la cryoablation ou la radiofréquence pour les lésions pulmonaires ou osseuses. Les événements indésirables ont été évalués selon les critères CTCAE. Les données radiologiques, y compris les images de CT et d'IRM avant l'ablation thermique et lors du suivi, ont été examinées par un observateur indépendant. Des évaluations cliniques, biologiques et d'imagerie de routine ont été effectuées. Le suivi comprenait une réévaluation radiologique à un mois, deux mois, puis tous les trois mois en collaboration avec l'équipe oncologique.

Les profils moléculaires des tumeurs ont été évalués conformément aux recommandations pour le cancer colorectal métastatique dans les milieux de soins. En résumé, les tumeurs sont analysées à l'aide d'un panel NGS ciblé de 29 gènes, incluant KRAS, NRAS, BRAF, TP53, PIK3CA, CTNNB1, et ERBB2, les plus fréquemment mutés dans le cancer colorectal. Les profils moléculaires des tumeurs étaient connus pour 80 % des patients référés pour une ablation thermique, et inconnus pour 20 % des patients car ils avaient été référés par un autre centre.

L'ADNc a été extrait en utilisant le Kit Maxwell[®] RSC ccfDNA (Promega), quantifié par le Kit

Qubit dsDNA HS Assay. Le cfDNA a été analysé en utilisant le panneau AmpliSeq™ Colon Lung Cancer (ThermoFisher) et les bibliothèques ont été générées automatiquement sur l'ionchef (Kit AmpliSeq™ for Chef DL8) à partir de 15 µL de cfDNA. Le séquençage de nouvelle génération (NGS) a été réalisé avec l'Ion Proton avec une profondeur minimale de 3000X. Les données de séquençage ont été analysées en utilisant la méthode BPER telle que décrite précédemment. Pour les patients sans données disponibles sur les tissus tumoraux ou avec des tumeurs de type sauvage, les échantillons plasmatiques négatifs ont été évalués par méthylation WIF et NPY-ddPCR (met-ddPCR) pour valider l'absence d'ADNc comme décrit précédemment.

Le ctDNA était positif pour 33 % des échantillons collectés au cours des premières 24 heures. La valeur pronostique du ctDNA de base a été analysée comme suit. Les patients ont été définis comme ctDNA positifs si du ctDNA était détecté avant l'intervention. Le ratio des patients ctDNA positifs/négatifs au départ était de 27 % / 73 %. La concentration moyenne totale d'ADN plasmatique avant l'intervention était de $15,53 \pm 19,00$ ng/mL, et le ratio moyen de mutation de la mutation prédominante, lorsqu'elle était détectée, était de $4,70 \pm 6,37$. Jusqu'à trois mutations d'ADN ont été détectées simultanément dans les échantillons. La médiane de survie sans récurrence était respectivement de 46 jours (IC 95 % : 29 - non atteinte) pour la population ctDNA+ et non atteinte (IC 95 % : 51 - non atteinte) pour la population ctDNA- ($p = 0,038$). Le rapport de risque pour la progression en fonction de la présence d'ADN tumoral circulant de base a été estimé à 0,14 (IC 95 % : 0,03 - 0,65, $p = 0,019$).

Puis, nous avons proposé de modéliser les trajectoires de deux populations en nous basant sur l'adaptation des modèles à effets mixtes pour les ensembles de données longitudinales avec ruptures temporelles inconnues. Nous définissons une structure hiérarchique qui découple la trajectoire moyenne de la population de la variabilité individuelle. En particulier, nous nous intéressons aux ruptures structurelles au sein des dynamiques de la population qui pourraient correspondre à une réponse au traitement ou à une progression infra-clinique de la maladie. La trajectoire de chaque individu est modélisée comme une trajectoire affine par morceaux. Cette trajectoire est considérée comme une courbe linéaire par morceaux d-dimensionnelle continue, ce qui permet de gérer les ruptures structurelles dans la progression de la maladie ou les changements de traitement. Dans notre cas, la dimension est de 1 et l'ensemble des temps d'observation est composé des temps 0, 1, 2 et 3, où l'ablation thermique a été réalisée entre les temps 0 et 1 avant la libération théorique de ctDNA.

Ce cadre est particulièrement adapté au contexte clinique où les ruptures structurelles (récurrence après ablation thermique) ne sont pas connues. Une trajectoire de population est entièrement définie par la liste des temps de rupture et la liste des valeurs de trajectoire à chaque rupture. La trajectoire est supposée affine entre les ruptures. Chaque individu est défini par la manière dont sa trajectoire s'écarte de la moyenne. Les paramètres du modèle sont calculés en utilisant l'estimateur du Maximum A posteriori (MAP), qui a été calculé avec l'algorithme MCMC-SAEM. Pour des raisons de faisabilité, nous utilisons les données synthétiques précédemment générées comme si nous avions recruté 200 patients, avec des étiquettes simples comme « progressif » (mauvais résultat) et « non-progressif » (bon résultat).

Dans le cadre de notre étude clinique, nous avons entrepris une double piste de recherche visant à nous aider à sélectionner les patients avant le traitement et à mieux comprendre les effets biologiques de l'ablation thermique. Le ctDNA s'est révélé efficace en tant que biomarqueur prédictif avant l'ablation thermique. Malheureusement, les quatre points temporels évalués n'étaient pas suffisants pour une compréhension approfondie des dynamiques.

Nous avons émis l'hypothèse d'une augmentation prononcée des niveaux de ctDNA après le

traitement en raison de la destruction des cellules par nécrose, apoptose et formation d'exosomes, ainsi que de la réaction immunitaire déclenchée. Nous avons supposé que les patients qui ne répondraient pas à long terme étaient ceux avec une maladie non diagnostiquée et non traitée, c'est-à-dire ceux qui auraient une forte baisse du ctDNA avant une forte augmentation. En parallèle, les patients ayant de bons résultats auraient une stabilité dans le faible niveau de ctDNA résiduel tout au long du temps.

Les deux trajectoires hypothétiques ont été produites par l'algorithme MCMC-SAEM, avec cependant quelques spécificités. Les trajectoires étaient en accord avec nos attentes, avec une forte augmentation observée chez les patients avec des résultats défavorables et un maintien constant de faibles niveaux chez les répondeurs, et ont même montré une légère hausse après l'ablation thermique. De plus, la détection des ruptures temporelles, notamment autour des points temporels 300/400, présente un résultat intrigant pour une enquête plus approfondie sur les dynamiques de la progression du cancer non traité. Cette rupture temporelle, environ deux semaines après le traitement, suggère un changement dans les dynamiques qui est concordant avec les évolutions connues du cancer colorectal. Deux semaines est typiquement le délai entre les lignes de chimiothérapie et les thérapies ciblées, et c'est aussi le délai pour interrompre les traitements dans le cadre péri-opératoire. Que ce résultat soit artefactuel ou non pourrait faire l'objet d'une étude prospective sur une cohorte plus large de patients traités par ablation thermique.

En marquant une rupture à l'Heure 400, le modèle fournit un aperçu important de ce que nous pourrions attendre des dynamiques de l'ADN tumoral par rapport aux phénomènes biologiques sous-jacents induits par l'ablation thermique. Utiliser un GAN sur 20 points en 4 dimensions semble être une approche plutôt compliquée. Les GAN sont connus pour être instables pendant le processus d'entraînement, en raison de la nature antagoniste de l'entraînement où les deux réseaux sont en compétition l'un contre l'autre. Comme ces réseaux sont généralement entraînés sur des données massives, notre GAN pourrait ne pas converger vers une trajectoire plausible de population pour les patients avec un bon ou un mauvais résultat. Il peut imiter les propriétés statistiques de notre petite base d'entraînement basée sur l'étude clinique menée pour ce travail, sans capturer les relations biologiques complexes entre ctDNA, ablation thermique et récurrence de la maladie.

4 Partie IV. Représentation planaire de l'arborescence vasculaire 3D

La détection et la segmentation des cathéters dans l'image sont des prérequis à l'autonomie au bloc. La présente étude a été approuvée par le comité d'éthique (IRB #00011928). Tous les patients consécutifs ayant subi une embolisation de l'artère utérine (EAU) entre janvier 2022 et décembre 2022 ont été examinés. Les critères d'inclusion comprenaient toutes les femmes ayant subi une EAU dans notre centre tertiaire. Les patients ont été exclus s'ils refusaient de participer ou s'ils étaient âgés de moins de 18 ans ou sous tutelle.

Toutes les procédures ont été réalisées par des radiologues expérimentés du service de radiologie interventionnelle, et les images radioscopiques en direct ont été enregistrées prospectivement. Les paramètres standard ont été utilisés sans modifications spécifiques du traitement dues à l'étude. L'EAU a été réalisée soit pour une hémorragie post-partum, soit pour une adénomyose, soit pour des fibromes utérins symptomatiques. Après la procédure, les patientes ont été hospitalisées pour la gestion de la douleur et ont été sorties le lendemain avec des prescriptions de médicaments antidouleur. Toutes les données DICOM ont été entièrement anonymisées, et les images ont été

récupérées et annotées par trois enquêteurs indépendants à l'aide du logiciel GIMP.

Un total de 152 patientes a été inclus. Pour chaque patiente, une image unique sélectionnée au hasard provenant de la vidéo scopique a été récupérée, contenant la pointe du cathéter et/ou le microcathéter au point d'injection. L'image entière a été incluse pour une analyse ultérieure. Les images individuelles ont été extraites des métadonnées DICOM et converties en fichiers en niveaux de gris standard pour faciliter l'analyse et le traitement. Les images en niveaux de gris contiennent des valeurs de pixels allant de 0 (noir) à 255 (blanc), représentant l'intensité de l'image à chaque emplacement de pixel. Les tailles des images varient de 512x512 à 1024x1024.

Le jeu de données a été divisé en un ensemble d'entraînement (période 1) de 112 images, un ensemble de validation de 20 images (période 2) et un jeu de données de test indépendant de 20 images (période 3) pour l'analyse des performances.

Les masques de segmentation correspondants ont été annotés manuellement pour chaque image et représentent une image binaire avec les pixels de fond en noir (0) et le masque du cathéter/microcathéter en blanc (1). En raison des limitations des données et afin d'éviter le surapprentissage et de rendre le modèle plus robuste, nous avons effectué plusieurs techniques d'augmentation de données. Les transformations incluaient : retournement horizontal aléatoire, perspective aléatoire, rotation aléatoire, contraste de luminosité aléatoire, bruit gaussien et recadrage central. Les algorithmes d'augmentation des données appliqués ont fourni plus de diversité aux ensembles de données, donnant 448 images dans l'ensemble d'entraînement et 80 images dans l'ensemble de validation. Le résultat de l'augmentation des données est représenté dans la figure ci-dessous. Un exemple est fourni dans la figure ci-dessous.

Nous effectuons l'évaluation de notre modèle sur un ensemble de test distinct composé de 20 images. Nous utilisons l'IoU (intersection sur union) et le score cIDice comme métriques d'évaluation. Les résultats de la segmentation sont représentés dans le tableau ci-dessous. Cette section est encore en cours d'investigation. Les performances sont fournies pour l'IoU (intersection sur union), et les résultats du cIDICE seront disponibles prochainement, car il s'agit d'un projet en cours.

Puis nous nous sommes intéressés à l'anatomie 3D de l'arbre vasculaire. Les techniques d'imagerie actuelles pour la vasculature comprennent des angiographies substractives en trois dimensions sélectives et mini-invasives à haute résolution. Ces images en trois dimensions, rehaussées par contraste, sont obtenues avant le traitement de divers troubles vasculaires tels que les anévrismes, les malformations, les fistules, les AVC et les sténoses. Elles fournissent des données essentielles pour la planification du traitement, notamment en ce qui concerne la voie d'accès vasculaire vers le trouble cible.

La navigation endovasculaire repose sur l'insertion de plusieurs dispositifs coaxiaux, y compris des guides, des micro-guides, des gaines, des cathéters et des micro-cathéters. Tous ces dispositifs présentent des caractéristiques spécifiques en termes de couple, de support et de poussée, qui détermineront les difficultés d'accès aux points cibles. Ces défis sont intrinsèquement corrélés à la courbure et à la longueur du trajet vasculaire, ainsi qu'à la friction pariétale le long du chemin.

Malheureusement, le rendu 3D de l'anatomie du patient ne fournit pas une représentation compréhensible des voies endovasculaires possibles et des défis associés. C'est pourquoi nous visons à fournir une représentation planaire simple et arborescente de l'ensemble de la vasculature cérébrale, qui peut aider les radiologues interventionnels à évaluer les obstacles avant le traitement endovasculaire, en utilisant une seule vue de l'anatomie du patient. Notre travail est également motivé par des études anatomiques pionnières de la vasculature cérébrale, réalisées via des dissections physiques de dizaines de cerveaux humains, que nous souhaitons maintenant automatiser de manière mini-invasive.

À partir d'une angiographie 3D, nous avons l'intention de créer une représentation planaire de la structure vasculaire 3D qui :

- Préserve la véritable longueur et le diamètre des segments vasculaires.
- Approxime les angles et les départs des branches artérielles.
- Estime la tortuosité des segments artériels par la courbure.
- Estime le couple accumulé du dispositif le long du trajet.
- Optimise l'utilisation de l'espace de représentation 2D disponible de manière efficace.
- Assure qu'il n'y a pas de chevauchement dans la représentation bidimensionnelle (2D) des artères.

Nous notons que la création de cartes des vaisseaux sanguins est un domaine de recherche actif. Ce qui distingue notre travail est l'attention portée aux longueurs et aux courbures des vaisseaux, souvent négligées par d'autres méthodes mais jouant un rôle crucial dans les interventions endovasculaires.

Notre population d'étude se compose de quinze patients consécutifs suivis pour des anévrismes intracrâniens non rompus dans un centre tertiaire unique. Les procédures ont été réalisées sous anesthésie générale. La première étape du processus consiste à positionner le système d'angiographie en forme de C autour de la tête du patient. Cet équipement se compose d'un bras en forme de C avec une source de rayons X et un détecteur qui peuvent pivoter autour du patient pour capturer des images sous plusieurs angles (Cone Beam CT). La première acquisition de tomographie par faisceau conique (CBCT) est effectuée sans injection de produit de contraste. Cette acquisition initiale sert à créer une image de référence ou une "masque" de la tête du patient et des structures intracrâniennes. Après l'acquisition de l'image de masque, un agent de contraste est injecté dans l'artère carotidienne ou vertébrale, permettant une visualisation améliorée de l'anatomie vasculaire. Une fois les deux acquisitions CBCT terminées, les données d'imagerie sont traitées pour générer une reconstruction tridimensionnelle (3D) de la vasculature intracrânienne après soustraction de la première acquisition de la seconde.

L'étude présente un pipeline complet pour construire une représentation planaire préservant la courbure de cette arborescence 3D. Un défi majeur des algorithmes de navigation des vaisseaux est d'éviter les artefacts topologiques induits par les limitations de la résolution des voxels. Lorsque les artères se chevauchent dans l'image, de grands voxels peuvent ne pas permettre une distinction précise entre elles, entraînant des "artères sautantes". Cela se produit lorsque le modèle interprète incorrectement les artères chevauchantes ou étroitement positionnées comme un seul vaisseau, ce qui fait que la représentation du vaisseau semble sauter ou se décaler. C'est pourquoi nous avons choisi de baser notre modèle sur des techniques d'imagerie invasive à haute résolution, telles que les angiographies, plutôt que sur des images non invasives en coupe transversale comme l'IRM ou les tomodensitogrammes.

Afin de produire une visualisation 2D sans chevauchement du réseau artériel cérébral tout en préservant les longueurs des vaisseaux, nous supposons que notre graphique est un arbre, sans cycles. Cela est cohérent avec les connaissances courantes sur l'anatomie cérébrale : à l'exception du cercle de Willis, un réseau artériel cérébral normal devrait être exempt de cycles. Malheureusement, cette forte hypothèse topologique ne se vérifie pas en pratique. En raison des limites de la résolution

des voxels, certains vaisseaux qui sont proches les uns des autres peuvent être fusionnés par notre traitement : notre traitement hallucine des "raccourcis" et des cycles topologiques qui ne correspondent pas à des chemins réalistes pour les interventions endovasculaires. Des anatomies patient anormales ou des ruptures de vaisseaux peuvent également créer des cycles qui doivent être pris en compte, mais qui pourraient être identifiés comme des arêtes "anormales" explicites au-dessus d'une structure normale d'arbre.

Un certain débruitage topologique est donc nécessaire. À cette fin, et pour des raisons d'efficacité computationnelle, nous fusionnons d'abord tous les nœuds ayant deux voisins afin d'obtenir un graphique plus léger. Cela réduit le nombre de nœuds de plus de 25 000 à seulement 2 000 en moyenne, tout en préservant la même structure topologique. Ensuite, nous créons un nœud racine qui représente le cœur pompant et le connectons à tous les nœuds situés au bas du volume. Enfin, nous assignons un coût à chaque arête orientée du graphique, ce poids encode une hypothèse sensible sur les arbres artériels, à savoir que le sang a tendance à circuler des vaisseaux larges vers les vaisseaux étroits. En utilisant l'algorithme d'Edmonds fourni par la méthode, nous trouvons alors un sous-graphe sans cycle du réseau vasculaire qui connecte tous les nœuds du réseau simplifié avec un poids total minimal, c'est-à-dire avec un flux "amont" minimal allant des vaisseaux fins aux vaisseaux épais. Nous affichons le résultat de cet élagage topologique sur le graphique haute résolution. Les arêtes induisant des cycles qui ont été éliminées par l'algorithme d'Edmonds sont affichées en vert, tandis que les autres sont colorées en fonction de leurs distances géodésiques au nœud racine.

Notre méthode est capable de capturer correctement la plupart du réseau artériel cérébral. Cela inclut des régions difficiles de l'anatomie cérébrale telles que les deux artères cérébrales antérieures qui courent parallèlement l'une à l'autre, entre les hémisphères gauche et droit, et sont correctement modélisées en tant que courbes distinctes. En pratique, nous constatons que la structure topologique des principales artères (avec un rayon supérieur à 2 mm) est bien préservée.

En revanche, deux principales sources d'erreurs persistent : • Certains branches de l'arbre peuvent être rejetées, surtout si elles sont fines et connectées au reste du réseau par un tournant aigu qui n'est pas très contrasté dans le volume 3D brut, • Bien que notre méthode gère efficacement les vaisseaux parallèles, elle ne gère pas bien les "croisements". C'est une conséquence du fait que notre procédure de nettoyage topologique ne peut supprimer que des arêtes du graphique et ne peut pas diviser un nœud de "carrefour" avec 4 voisins en deux courbes séparées.

Heureusement, ces artefacts de segmentation n'affectent qu'une infime portion du réseau artériel cérébral. Les erreurs se produisent généralement dans des régions trop éloignées pour être accessibles avec les dispositifs endovasculaires actuels. Néanmoins, obtenir une topologie de réseau fiable et précise à partir de nos images 3D est l'un de nos principaux objectifs pour les travaux futurs.

Pour des raisons de visualisation, nous enrichissons notre graphique avec des caractéristiques de nœuds qui sont calculées de manière récursive : • La distance géodésique à la racine, • Le volume des vaisseaux en aval, et calculé en identifiant chaque arête avec un cône tronqué. • Les coordonnées XYZ du barycentre des vaisseaux en aval, que nous affichons directement sous forme de couleurs RGB.

Au cours des dernières décennies, de nombreux algorithmes ont été proposés pour dessiner des graphes sur un plan 2D. Les méthodes standards telles que l'algorithme Kamada-Kawai nous obligent à choisir entre préserver les longueurs de chemin ou fournir des visualisations sans intersections. Bien que des progrès significatifs aient récemment été réalisés dans le démêlage des nœuds et courbes en 3D, nous avons constaté que les méthodes de dessin basées sur les forces existantes ne pouvaient pas afficher de manière fiable en 2D les topologies complexes des réseaux artériels cérébraux.

Afin de satisfaire toutes nos contraintes, nous proposons donc une méthode originale pour la représentation planaire des graphes, spécialement adaptée aux besoins des radiologues interventionnels. Nous visons à produire des dessins esthétiquement plaisants de la neurovasculature qui améliorent la lisibilité du graphe, sans compromettre les considérations anatomiques.

En premier lieu, nous calculons une représentation planaire de l'arbre en utilisant un algorithme récursif, puis nous courbons les vaisseaux pour mieux refléter leur courbure locale. À cet effet, nous ajustons d'abord un plan à chaque point du réseau artériel 3D en effectuant une Analyse en Composantes Principales sur des quartiers géodésiques. En partant de la racine, nous orientons ces plans de manière cohérente et calculons, pour chaque paire de bords 3D consécutifs, un angle de courbure orienté dans le plan local 2D. Cette courbure correspond à une courbure cible pour nos encodages 2D. Enfin, l'opérateur de lissage est nécessaire pour éviter trop de dérive dans la formule de courbure. Pour équilibrer cette nouvelle force de courbure avec la répulsion électrostatique, nous itérons simplement jusqu'à convergence la mise à jour suivante sur les angles des bords de l'encodage 2D.

Titre : Modèles mathématiques pour la radiologie interventionnelle personnalisée : Application au traitement du cancer

Mots clés : Modèles statistiques génératifs, détection d'anomalies, déformation difféomorphique

Résumé : Les méthodes statistiques contemporaines ont le potentiel de changer la pratique médicale au sein même de nos blocs de radiologie interventionnelle. Ce travail pose quelques jalons pour l'avenir des interventions autonomes dans notre domaine spécifique avec pour fil conducteur la prise en charge des patients atteints de cancer.

Nous explorons d'abord l'impact transformateur de l'IA sur les capacités physiques des radiologues interventionnels. À terme, l'autonomie au bloc passera par la robotisation et l'analyse d'images produites en direct. L'intégration de la vision au geste permettra d'avancer vers ce bloc autonome.

Nous proposons ensuite une méthode de détection d'anomalies ne dépendant pas de l'organe exploré à partir d'une imagerie en coupe. En exploitant le cadre du Large Diffeomorphic Deformation Metric Mapping (LDDMM), cette approche met en valeur une reconstruction d'objet améliorée et une segmentation de la lésion. Dans le même cadre, nous proposons un classificateur,

permettant d'établir une nouvelle lecture de la sélection des patients avant l'intervention.

Afin d'aller au-delà des images, des données cliniques issues de l'analyse de l'ADN tumoral circulant sont intégrées dans une étude prospective spécifiquement menée pour ce travail. Des réseaux antagonistes génératifs (GAN) et des algorithmes d'approximation stochastique (MCMC-SAEM) sont utilisés pour prédire les trajectoires des patients. Cette approche permet d'explorer de nouvelles trajectoires, améliorant notre compréhension de la progression de la maladie et de la réponse au traitement en relation avec l'ADN tumoral circulant.

Enfin, nous explorons des techniques de visualisation avancée pour l'anatomie vasculaire 3D in vivo et ex vivo. Nous proposons une représentation planaire de l'anatomie, offrant des possibilités en termes de navigation et de classification de l'arbre vasculaire.

Ensemble, ces sections offrent des solutions à certains défis menant vers la réalisation d'interventions autonomes guidées par l'image.

Title : Mathematical Models for Personalized Interventional Radiology : Application to Cancer Treatment

Keywords : Generative statistical model, Diffeomorphic deformation, sparse lesion detection, organ agnostic detection

Abstract : The integration of computer vision into Image-Guided interventions has the potential to change our medical practice. This work lays some bricks for the future of autonomous interventions in our specific field regarding cancer patients, addressing key components necessary for its realization. We first explore the transformative impact of AI on the physical abilities of interventional radiologists. We emphasize the need to navigate technical and ethical challenges. Interdisciplinary collaboration and robust evaluation processes are highlighted as essential for the safe integration of AI into clinical practice. We then propose an organ agnostic method for detecting focal anomalies on volumetric cross-sectional imaging. Leveraging the Large Diffeomorphic Deformation Metric Mapping (LDDMM) framework, this approach showcases enhanced object reconstruction and precise lesion localization. In the same framework, we propose a classifier, where

patient selection presents unique challenges due to the complex benefit/risk ratios. To go beyond images, clinical data from tumor DNA analysis is integrated into a prospective study specifically conducted for this work. Generative Adversarial Networks (GAN) and Modelling Atlases Using the Markov Chain Monte Carlo - Stochastic Approximation Expectation-Maximization (MCMC-SAEM) Algorithms are used to predict patient trajectories. This approach enables the exploration of new trajectories, enhancing our understanding of disease progression and treatment response in relationship of circulating tumor DNA. Lastly, we explore advanced visualization techniques for in vivo and ex vivo 3D vasculature. We propose a planar representation of undescribed anatomy, offering a promising avenue for further exploration and understanding. Together, these sections offer solutions to parts of the realization of autonomous interventions in our field.