

Cancerworld

AI in Surgery

Adriana Albini / 8 March 2023



The fourth online seminar of the SPCC Educational Project on Artificial Intelligence in Cancer Care, took place on 25th January 2023 and was moderated by **Pietro Valdastri**, Professor and Chair in Robotics and Autonomous Systems, University of Leeds, UK.

Developing Robotic Surgery: An Overview on the Latest Innovative Technologies

Elena De Momi, Professor in the Department of Electronics, Information and Bioengineering at Politecnico di Milano, Italy, began her presentation with a brief history of robotics for surgery. The very first robot used in a surgical procedure was the Unimation Puma 200, an industrial robot, which performed a stereotactic brain biopsy in California in 1985. In the 1990s, the Robodoc Surgical Device was used for orthopaedic surgery, while the Probot system carried out transurethral prostatectomy. Research into the application of remote surgery was first spurred by space exploration and then translated to the military, with a view to improve surgical capacities on the battlefield. SRI International and DARPA (Defence Advanced Research Projects Agency) devised a forefront and backward system: the Medical Forward Advanced Surgical Treatment (MEDFAST) was an armoured vehicle containing a robotic arm that could be driven to the battlefield. The Mobile Advanced Surgical Hospital (MASH) was a console that allowed the surgeon to teleoperate the robotic arm at a safe distance from the MEDFAST. In 1995, SRI created a spin-off company, Intuitive Surgical, Inc., which took their minimally invasive robotic surgery system called **Da Vinci** (in honour of Leonardo) into hospitals worldwide. The system performed its first human surgery, a cholecystectomy, in Belgium in 1998, and in 2000 received approval by the FDA.

There has been a number of generations of Da Vincis: S, Si, Xi, X and the latest concept, the Da Vinci SP (single port). Vision plays a crucial role in surgical robotics, so the image capture components have increased in resolution and precision over the years, and the system's ergonomics have also

improved with time. In the Da Vinci X, all the arms can play different roles. For instance, the endoscope can be attached to any arm. The latest product from the Da Vinci family is the single port, that allows the surgeon to control three instruments and an endoscope through a single incision. While in the literature we can find plenty of clinical trials related to the previous versions of the Da Vinci, there are still few that demonstrate the clinical effectiveness of the Da Vinci Single Port (see Ryan Dobbs et al., Single-port robotic surgery: the next generation of minimally invasive urology, World Journal of Urology, 2020); the cost effectiveness is still to be proven, as well.

Apart from the concept of teleoperation, with the Master Tool Manipulators on the surgeon's side and the Patient Side Manipulators on the patient's side, the biggest advancement of the Da Vinci is the surgical instrument which provides the surgeon with greater dexterity when manipulating soft tissue inside the patient's body, thus guaranteeing safety and precision in the surgical execution. As already mentioned, we can find plenty of clinical trials and studies describing the clinical application of the Da Vinci in different surgical fields: **cardiothoracic** (cardiac robotic lobectomy); **general surgery** (rectal cancer resection, gastrectomy); **head and neck surgery** (transoral surgery for oropharyngeal carcinoma, laryngeal tumours, parapharyngeal space tumours, thyroidectomy); **urology and gynaecology** (prostate cancer, partial nephrectomy, complex benign hysterectomy, cervical and ovarian cancer).



Patient Side Manipulators: robotic arms teleoperated by the Master Tool Manipulators, they mount the surgical tools.

Endoscopic Camera Manipulator: robotic arm that is also teleoperated by the Master Tool Manipulators, it holds the endoscope.

Why is image quality so important? Because artificial intelligence takes advantage of the amount of information provided by the image. With endoscopic images we can develop machine learning algorithms to provide, for example, automatic detection of regions within the images; segmentation, whereby the organ or the tissue area is not only detected, but also contoured. When detected, it can also be classified in terms of the histology, let's say, a virtual biopsy. We can also perform 3D reconstruction when we have three images, for example in the case of the Da Vinci systems. The spectrum of possible smart applications can thus be expanded. For example, we can recognise automatically in which part of the surgical workflow we are, what task the surgeon is performing, for example, suturing. We can establish safety regions, such as a vessel or a nerve which must be preserved and not touched by the robotic instrument. We can standardise the procedure, make a generalisable knowledge about how to perform the surgery. This, of course, can be beneficial for training and skill assessment of young surgeons.

The Da Vinci community has been very active in research, because Intuitive Surgical has provided an open system to university research centres around the world, so researchers have been working on medical robotics now for 10 years (see Pierre Dupont et al., A decade retrospective of medical robotics research from 2010 to 2020, Science Robotics 2021). **Where are we now?** Da Vinci has been a major player in robotic laparoscopy for the past 20 years, because Intuitive Surgical held all the patents related to these applications, but some patents are now expired, and many competitors and alternatives are emerging in the market. Focussing on robotic laparoscopy, there are 15 systems that have gained some form of clearance, from the FDA or other agencies. Among them are the Senhance Surgical Robotic System, first developed by the Italian Sofar SpA, then acquired by Asensus in the US; the Versius system from CMR, the Symani Surgical System, from another Italian company, Medical Microinstruments; and the Hugo RAS system by Medtronic, which is possibly the most credible alternative to the Da Vinci. The TRL (technology readiness level) of these systems is 9+, and five of them have had sales in more than one country. Yet, according to the available literature, only four are associated to some form of clinical trials.

Looking at numbers, there is still room for improvement and many possible new applications may emerge. 15% of all general surgery procedures in 2020 was performed robotically (from 1.8% in 2012); 1.5 million procedures were performed using the Da Vinci surgical system (with a 15% annual growth); the amount of money which has been invested in this field of application has also increased. The total market value of surgical robotics was estimated at \$8.3B in 2020, to reach \$33.6B in 2026.

Among other important innovations in the field are **continuum robotics**. Continuum robots have almost infinite degrees of movement. Their level of flexibility allows them to follow non-linear trajectories, such as lumina. In research, they have been applied to different surgical locations, for example, neurosurgery, abdominal surgery, cardiac surgery, and urology. On the market only two systems are present at the moment, both for lung biopsies, the Monarch by J&J, and the Ion by Intuitive Surgical, the same developer of the Da Vinci system. There is still very little literature published on the results of the first patients. As to the future, according to the already mentioned review paper by Dupont, the level of autonomy of robotic systems from the human operator is going to increase, widening the spectrum of possible applications. **What are the novel trends?** Novel trends are associated to novel needs. So, the 5G communication protocol is making real the possibility of having wireless connection and remote surgery, remote mentoring, remote training, or remote diagnosis, which can be performed by an expert located far away from the patient. Autonomy is increasing thanks to AI. We will need to train a new generation of surgeons who understand how to use the new system, but without forgetting the manual technique. Because if the system fails, they have to be able to go back to manual or open surgery procedures. As to regulatory issues, new medical device regulations are being brought in, for instance in Europe, and artificial intelligence needs to be robust and trustworthy in order to enter the market and get the CE mark.

Prof. De Momi concluded her presentation with a slide about some of the research lines that are being conducted at the Politecnico: processing of endoscopic images for 3D reconstruction; autonomous control of the camera motion; and application of continuum robots in the neurosurgical field.

AI Visualization Systems for Next Generation Treatments


Danail Stoyanov is Professor of Robot Vision at the Department of Computer Science and Director of the Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS), UCL, London, UK. There are around 4 million procedures in the UK performed each year, and probably a few hundred million worldwide. The systems that perform these procedures are increasing in complexity. We have robotic systems, such as dual console robotic theatres; there is much connectivity between the OR, or the interventional radiology suite, and hospital data management systems; and many

different devices are available. **The general trend of these new tools is towards being less invasive.** To achieve this, generally speaking, we need to have a visualisation of the surgical site, and that visualisation may come from a variety of different types of scopes. We now have stereo cameras, monocular cameras, flexible, rigid, etc. We use energies that are beyond the visible spectrum, so fluoroscopy or fluorescence or different types of imaging modalities. There is a plethora of imaging systems that provide a visualisation of the surgical site, to aid the surgeon. Laparoscopic cholecystectomy was a driver for minimally invasive surgery, but many other procedures, including trauma procedures, have transitioned to being minimally invasive over the past 15-20 years. If we look at an emergency case, like an appendectomy, 20 years ago probably around 30-40% of appendectomies were performed using a minimally invasive approach. Now the number is over 90%. And this trend is very likely to continue. Because we are using cameras, and most cameras today are digital, a huge amount of data is generated from the human anatomy and from the different surgical sites. As the camera resolution becomes higher and the different energies that we acquire become more variable, a lot of data is created that we can potentially utilise, because it contains information about the way that organs and tissues look, information about what happens during surgery, about what instruments are used, how they are used, and how they interact with the patient. This vast amount of data is currently not utilised very effectively, and we need to find solutions, because storing it in perpetuity is difficult. An hour of surgery in HD, even when compressed, is probably a couple of gigabytes of storage. So, we need to have efficient ways to analyse which parts of the data stream are important, what we want to do with them, how we can quantify them and what type of assistive systems we can build upon them. In the human visual system, parts of our vision can detect objects, part can detect motion, other parts may understand 3D space. We can replicate some of these functions for clinically assistive systems. Prof. Stoyanov showed two short videos of examples of AI assistive systems, the first, used in gastroenterology, can detect polyps or parts of the tissue that require more detailed investigation. The second, used in surgery, is able to differentiate between a duct and an artery: these are critical structures that the surgeon should manage appropriately in order to reduce the risk for the patient.

LapChole Critical Structure ID

ML performance in challenging scenarios

Successes?



Concerns over opacity/obstruction of anatomy but mitigated by ability to turn on/off




Flickering of artery/uncertainty of algorithm is fine as it mimics human thinking




Would depend on use case – for a verification tool it would be fine

Borderline cases?



Clear and realistic detection



Anterior view is good for verification and mitigates errors seen in posterior view



Sensitivity>Specificity is better – good to err on the side of caution

Systems like these have been in development for quite a while, but only in the past five years they have become robust enough to be used clinically, and are now entering clinical practice as products.

One of them is **CADDIE**, developed by Odin. This is an AI medical device for the detection and characterisation of polyps. It will be part of the Olympus Endoscopy business, and it is only used in the European market at the moment. Over a thousand patients have already undergone treatment with CADDIE, and the number is likely to grow dramatically over the next few years. Even a fairly simple system like this can increase adenoma detection by 5, 6, 7%, depending on clinical factors. Systems like this also provide a mechanism to generate the clinical report or to get a quantitative view of what occurs during all endoscopies in a given clinical unit. We may be able to visualise which parts of the colon were observed, collect images from those, understand the bowel preparation, understand how many polyps were detected, how many were resected, what size they were. So, we can have quite a detailed report of an endoscopy procedure, way beyond what is happening in practice today. A head of endoscopy could then log in through the Odin system and basically see a report of everything that happened in the hospital theatres on that day. Another example is the Medtronic technology to detect critical structures in laparoscopic cholecystectomy (lap chole). This is not as yet a medical device, but it will be in the near future. In surgery, we do not normally detect the structures for reporting purposes, but it would be valuable information. It is likely that systems like this are going to emerge as adding an extra level of safety on top of the existing practice. Quite a lot of generalisation still needs to occur for them to be used in clinical practice. They must be able to handle different levels of variability, with different cameras or instrumentation being used, etc., but eventually they will be part of the next generation of robotic systems. Being able to run algorithms like this allows for a detailed record of what happens daily during surgery. It could be useful as continuous feedback, for monitoring trainees, for generating operative notes, and so on. A number of potential clinical uses could create value out of the video stream, which is currently underutilised.

One of the key issues we face as a community is the availability of data. Although a lot of surgery today is done using digital cameras, the data is often not stored. And even when it is, it is not stored in a structured way, it is not made available generally for research or development. This creates a major bottleneck in deploying the new systems to different specialties, and testing their robustness, their generalisation. The ability to utilise the data is still a big challenge in the development of AI for surgical applications. The available data is growing, but very slowly. Taking, for example, a recent dataset that Prof. Stoyanov's group has been involved in: 50 videos from robotic assisted radical prostatectomy detecting instruments in certain phases. It is very small when compared to the size of data sets available for autonomous vehicles or other industrial uses, which are in the region of tens, hundreds of thousands, maybe even millions of sample size cases. Both clinical communities and engineering communities need to work on increasing the availability of data, because that will allow for the exploration of new potential AI-driven capabilities for different specialties. One way to avoid the need for data is, of course, to use simulation or unsupervised learning. For instance, we can simulate a CT of a colon in a virtual environment. We can estimate how cameras fly through that CT and we may be able to generate depth estimation systems, and depth estimation is one of the most important navigational cues. There are a lot of people working on this, but it is very challenging. Also, from a regulatory point of view, regulators are nowhere close to understanding how to treat simulated systems that are trained on simulation data.

As these AI systems become more robust, the device companies do not just have video data, but also device data. If we are using a robot, we have kinematics, we have system information data. If we are using energy devices and we have the energy profile of what the instruments are doing, linking that instrument data to the video data could be extremely powerful. Previously, if we had an energy instrument, we could not really understand what the energy profile meant, because we never had the context of what is happening inside the body. But now if we link the video to the energy instrumentation, we can actually understand how the energy instruments are working, and we can potentially use that information to perform risk mitigation or to optimise workflow or understand

system reports in robots, and that will lead to better critical systems. Prof. Stoyanov finished his presentation with a video of how AI enhanced visualisation could look like in the future. The AI system can detect where the instruments are within the field of view, what activity they are performing, how they are articulating. It can detect the instruments, and where certain instrument usage is occurring, but also when the instrument goes out of view, which is important to know, as we may damage tissue if the instruments are doing something outside the field of view. We can highlight particular structures or focus on particular parts while fading out others, and there can be many other functionalities for safe minimally invasive systems.

Methodological innovation and ethical aspects of AI in surgery: two sides of the same coin



Sara Moccia is Assistant Professor at the BioRobotics Institute and Department of Excellence in Robotics and AI, Scuola Superiore Sant'Anna, Pisa, Italy. Artificial intelligence, as we know, is now playing a major role in processing medical data, and surgery is no exception. This may be due to increased computational power, larger memory to store and process data, and larger annotated data sets to train algorithms. When today we talk about AI in surgery, we are referring to **weak AI**, algorithms that are able to solve or tackle specific and tailored challenges. And most of the time, when we talk about artificial intelligence in surgery, we talk about machine learning, which is a branch of artificial intelligence. It also includes deep learning, which is becoming more and more successful in addressing all the challenges, for example, relevant to intra-operative image analysis. So, we are far away from **general AI**, which is capable to perform operations like humans, and even more so from **super AI**, whose algorithms are **more capable** than humans. So, we are not really talking about machine consciousness at this stage. Weak AI today in surgery can be employed for tasks such as surgical risk analysis, preoperative planning, biometry estimation, image registration, video mosaicking, and follow-up prediction.

Have we reached the final goal? In March 2004 a workshop was held, called "OR2020: Operating Room of the Future" ([Cleary et al., 2004. OR2020 Workshop overview: operating room of the future. International Congress Series 1268](#)). It was attended by around 100 clinicians, engineers, and scientists to discuss clinical and technical requirements for the OR of the future. The chosen target date was 2020. When, at a follow-up workshop in 2019, participants were asked, "What has really changed when you are entering the OR of today as compared to the setting in 2004?" the general reply was that **they did not perceive any disruptive changes** ([Maier-Hein et al., Surgical data science –from concepts toward clinical translation, Medical Image Analysis 2022](#)). Most of the innovation today is associated to changes in the hardware, for example in the robotic platforms or the imaging devices to acquire intra-operative images. Yet, artificial intelligence is seen as a key enabling technique for the OR of the future. So, what is hampering it from entering the real surgical practice? Of course, there are methodological barriers, such as data availability, self-supervised learning, or semi-supervised learning. But there may also be some ethical issues arising from the use of an intraoperative AI decision support system. A major one is the **algorithm bias**. We should acquire and build data sets that are diverse and inclusive, so that they are representative of the populations that stand to benefit from these applications. How far are we from this? A framework for designing and deploying a machine learning algorithm should be as follows: we have data coming from sensors, imaging devices, robotic platforms, and electronic records. These data should be annotated by surgeons or by OR professionals. Once data and relevant annotations are available, then, the AI professional should train, develop, and deploy algorithms in a continuous framework, called Machine Learning Operations, (MLOps). So, we start from data and model, then we train and release our model, we assess its performance in the actual practice and, perhaps, go back to change the data, for example, increasing the size of the dataset or changing the model. But the truth is that we do not have enough data available. We have big data in terms of numbers, but we miss the annotation. There are some international initiatives, such as the MICCAI Grand Challenge, but the


data released by these challenges has limitations, associated to the fact that it is collected in a single surgical centre by a single surgeon, and is annotated by few experts. Also, most of the time there is no information about the patient population, such as gender, ethnicity, age, professional activity, level of education. **This is not just a technical problem.** So, the problem is not so much that we lack the tools to annotate the data, but that we have biases and inequalities that are already embedded in the healthcare system. When building our algorithms, we must avoid to sustain, perpetuate, and exacerbate existing biases and inequalities embedded in healthcare tools, therapeutics and systems that see certain groups receiving substandard or inappropriate care. For instance, we may have hospitals with more male than female surgeons, or with only white patients between the age of 40 and 60.



But in practice...



FLARE 2022



Sharing Progress in Cancer Care

Aside from bias and privacy issues linked to clinical data, little consideration has been given to ethical concerns that arise from the use of an intraoperative AI decision support system. We should be able to explain not just how but also why an AI and machine learning algorithm takes a decision. This, of course, can be addressed methodologically, but we have far to go to reach the optimum. Equally, we must remember that we are not explicitly training our algorithms to embed the moral requirement to do good and avoid harm. This is something that is transferred to the algorithm because we as human implicitly embed such moral requirements. Ultimately, the full moral responsibility lies with the surgeon, both if in agreement or in disagreement with the decision support system. In 2019, the European Commission presented Ethics Guidelines for Trustworthy Artificial Intelligence. According to these Guidelines, trustworthy AI should be:

1. lawful - respecting all applicable laws and regulations
2. ethical - respecting ethical principles and values
3. robust - both from a technical and social perspective

(<https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>). As for the ethical aspect, the European Commission identified four principles which are: respect for human autonomy, prevention of harm, fairness, and explicability. To help AI professionals, the European Commission also provides an assessment list to check if an artificial intelligence algorithm is trustworthy. The framework is open-ended, so it needs to be tailored to each specific field. So, how do we

operationalise this framework in surgery? First of all, we should separate the different phases of design, testing, and large-scale deployment. When designing artificial intelligence, ethical issues may arise both from the dataset and from the algorithm. We may have biases when we have insufficient representation of a certain group, age, gender, ethnicity and so on. As for the algorithm, we may have **infallibility bias**, when we believe our algorithm to be infallible; **cohort bias**, when the algorithm only works with traditional or easily measurable groups; and **label bias**, when it only works with labels annotated during training. Moving on to the use phase of a ML system, there might be issues arising from the interaction of the surgeon with the hardware used to acquire data. There may be, for example, **accidental mistakes** that arise from an interaction of the surgeon with the camera, such as changing its angle, or with the robotic platform, that may compromise the results of the AI algorithm; **automation bias** when the surgeon overrelies on the AI model; **feedback loop bias**, i.e., a cascade errors generated by previous ones, validating the algorithms on inaccurate evidence; and **dismissal bias**, a cascade of errors causing the entire dismissal of the system.

As for ethical issues arising from the large-scale deployment of an ML algorithm in surgery, we must consider fairness and sustainability, such as economic and environmental costs, disparity in private and public hospitals, mid- and high-income countries, etc. In terms of privacy, recording devices should be ward-privacy sensitive; we also need to have consent from patients for using their data for training or testing algorithms, but also for using the algorithms in the actual surgical practice. Last but not least, we need transparency. The regulation of AI systems is very strict, but most of the time the assessment is performed from the technical perspective, while the ethical audit is still something that is left out, while it would greatly contribute to reduce the opacity of the algorithms.

To move forward the adoption of AI in the operating room, we need more than just adding engineers, computer scientists and clinicians, we should also consider including other professionals such as, for example, artificial intelligence philosophers, to assess the impact of the algorithms, not just from a quantitative point of view, but also from the ethical perspective, the effect they can have on surgeons and patients.

The current landscape of AI for gastrointestinal cancer screening and detection

Keith Obstein is Associate Professor of Medicine and Associate Professor of Mechanical Engineering at Vanderbilt University, Nashville, TN, US. There are a number of GI cancer screening methods. **Endoscopy** is the primary screening tool for GI cancers. Upper endoscopy, or more precisely esophagogastroduodenoscopy (EGD), is a video camera inserted down the mouth, oesophagus, stomach and first part of the intestine. It is used for conditions like Barrett's oesophagus, oesophageal cancer, as well as ulcers, gastric cancer, etc. Lower endoscopy composes of colonoscopy, as well as **capsule colon endoscopy**. Then, we have enteroscopy, which looks at the small bowel, and we can use a long video camera called enteroscope or a colonoscope, or wireless capsule endoscopy, which is a capsule with a camera that takes pictures and transmits them to an external recorder. The difference between the capsules and the other interventions is that a capsule is purely diagnostic, meaning you can see, but, at this point in time, you cannot intervene using its modalities.

Colonoscopy is the only screening test that does not only detect a colon cancer, but can actually prevent it. Colon cancers typically form from polyps and adenomas that grow inside the colon, which eventually turn into cancer. With screening, if we find polyps or adenomas, pre-cancerous lesions, with a colonoscope and take them out with polypectomy, the risk of them becoming cancer becomes minimal. A very large, seminal study from the New England Journal showed that the higher the

adenoma detection rate, the lower the likelihood that a patient will develop an interval colon cancer, that is a colon cancer that either was missed or that developed in the interval between colonoscopy appointments. Vice versa, if the adenoma detection rate is low, the hazard ratio goes up significantly.

Another form of GI cancer screening is **imaging**, and there are different modalities. CT and MRI are beneficial for somebody with a hereditary pancreatic cancer. Imaging is good at surveying once a patient already has cancer. Sometimes it can happen that cancers are detected as incidentalomas, that is, they are found accidentally while looking at imaging for another reason. Now we can also do hormonal studies, meaning radionuclide scans or PET scans, and overlay them to CT images, thus obtaining a higher degree of lesion identification.

Then, of course, there is **lab screening**. Most of the times blood tests are not really used to detect the colon cancer, as much as they are used to follow it. We can find tumour markers in the blood, and we use those to assess response to therapy, surgery or medications, chemotherapy, radiation, and so on. Cancer can also be detected with urine cytology. If there is blood in the urine, one of the things to look for is a neoplasm; also, we can do stool tests for colon cancer screening, and there are several at-home testers. Fluids, like ascites, pleura, pancreatic can also be analysed for cancer.

Moving on to AI, Professor Obstein took some entertaining but effective examples to highlight the many complexities of computer-aided detection (CADe) and computer-aided diagnosis (CADx) that are everyday operations for a human brain. The first was from a football match: from detecting the players on the pitch, to diagnosing what team they belong to, to identifying the ball, the goal post, triangulating distances, and eventually, making some sort of strategical assessment. Once we have made a detection, we can proceed to make a diagnosis. Vision in AI is quite different from human vision, and that is another factor to keep in mind. As humans, we can tell quite clearly the difference between an image of a fried chicken leg and a poodle, or a blueberry muffin and a dog's face, or a bagel and a curled-up dog. But for a computer it is not so simple. The last example was what an autonomous driven vehicle sees while it drives down the road, all the data that is being entered, all the manipulations, calculations, traffic lights, where other cars are, where the road ends. A road is a constantly moving environment. Similarly, the GI tract is a constant dynamic setting, and as such it requires a highly complex series of algorithms, detection, and training sets. **How can AI decision support help?** Simply detecting a polyp is not enough, we then have to ask a number of questions: can I do an optical biopsy? What is the likelihood that this polyp will turn into a cancer? Do I even need to take it out? Hyperplastic polyps, for instance, will never turn into cancer, and so, to remove them only exposes the patient to risk, with no real benefit. In the case of inflammatory bowel disease, can I objectively assess the severity of somebody's Crohn's disease or ulcerative colitis and determine a therapeutic regimen, without being influenced by other factors, such as my relationship with the patient or my training experience? What about the quality of our exams? Some tests are not clear, if there is blood in the stomach for example. Do we need to repeat them? If so, when? And how efficient was the process? AI decision support can help with all of that, and could also potentially be capable to autonomously move an endoscope through the colon in a similar way to an autonomous vehicle. We have the power, we have the computer abilities, but can we have an endoscope "drive" autonomously through the colon? And, of course, AI could help to identify the patients who might need screening or benefit from surveillance, or are up for their next colonoscopy based on prior polyp data. AI can also write letters of approval to insurance companies or different health boards saying that the patient should have a certain test and why. This would save clinical staff a lot of time. To recap, we have imaging, we have navigation and ability to drive, we have diagnosis, management assistance, even assistance in patient facilitation, and letter writing efficiency. The field is wide open.

Artificial Intelligence (AI)

- Computer-Aided **Detection** (CADE)
- Computer-Aided **Diagnosis** (CADx)



Where are we now? Professor Obstein showed a video with a polyp detection algorithm. The corners of the screen change colour and give a percentage when a polyp has been detected, allowing the team to spend extra time looking in that area to find it. Other detection algorithms actually hone in on where they believe the polyps are, with boxes that flash over them. Unfortunately, during surgery, for instance to remove a polyp, the algorithm can sometimes interfere with the visualisation, so, the ability to shut it off at certain points to let the endoscopist assess the real margins or extent of resection is key. Looking at AI in quality and clinical management, there are pieces of software and algorithms for procedure documentation and reports, that can fill in answers to questions like, what was the bowel preparation? What instruments did I use? These algorithms and software can detect what we used, maybe what the polyp was, so a pathologic optical diagnosis is included in the report. This provides more extensive information and saves time to surgeons and nurses, who would otherwise have to write the reports themselves. Software can be employed to enrol patients in studies. Let's say we are trying to look at inflammatory bowel disease patients and grade them by level of inflammation, we want to put them forward for a new study drug, but we might not even know that such study is going on. We upload video and data to the Cloud, where a piece of visual detection software looks back at our data and assesses if the patient would benefit from enrolling in a particular study or from using a novel therapy or drug, and then connects the pharma company with the provider and the patient. This is a great tool, especially when dealing with rare diseases that might only be seen once or twice in a particular institution. Now we can have huge amounts of data from around the globe that we know are standardised and objective. As to autonomous manoeuvring, we can have a magnetic-assisted endoscope that allows us to drag a colonoscope through the colon as opposed to pushing it, and we can also use algorithms to detect the centre of luminal depth. With supervised "driving", we can manipulate the endoscope through the colon autonomously, locally, or even remotely. When doing biopsies, we can find polyps and lesions. We know where the forceps is going to come out, and it is possible to automatically triangulate that so that it comes right onto that polyp; we can get a stabilised biopsy of the tissue, and do the extraction. So, we can change the learning curve, we can change the dynamics of the environment, and also have intervention assistance. These are all very powerful tools.

What are the pros and cons of current platforms? Among the pros are: the potential to **improve**

low performers or low-quality screening exams with the detection of polyps; **finding the “unfound”**, for instance, we might be looking for Barrett’s oesophagus and maybe we do not detect an area of dysplasia; **real-time diagnosis**: not needing to take something out and expose the patient to the risk of removing a lesion would be hugely advantageous clinically, but also financially; speeding up and optimising the procedure **documentation**, also is incredibly useful. Among the cons, AI systems are only as good as the **data** they were trained on. If we show them something new or different, depending on the power of the algorithms used, we might not be as successful. Another point is that **AI cannot see what we do not show it**. So, let’s say, in endoscopy and colonoscopy, if I am not very good technically and cannot see the lumen or detect all the surfaces, and I see nothing at all, AI is not going to see it either. Because I did not see it, I did not point the camera for AI to see. The **interface** can be distracting, with its flashing boxes, etc. Sometimes it creates cognitive overload, but other times, it can actually block the margins of what we need to see. For instance, we can have lines on surgical tools, but really what we want to see is where we are making that cut. We all hope that these new tools are going to be an improvement, but studies on clinical practice that are coming out now indicate that, maybe, they are not as successful clinically as we would have liked. For instance, a study published in the American Journal of Gastroenterology in 2022, shows that AI for colon polyp detection actually reduced the high-performing endoscopists, those with really successful adenoma detection rates. They started using AI software, and their detection rate went down significantly. What happened? It could be a result of information overload: AI may be giving too much irrelevant information, that, when it shows something important, we might not pay attention. On the other hand, blindly trusting the algorithm could also cause harm to the patient.

AI in GI is advancing rapidly, in fact, GI is possibly at the forefront of this innovation. There are top quality HD videos and recording abilities, which are used clinically, but we need to understand the benefits and limitations of these platforms and systems. We need to seriously think of how they can be applied to accomplish the goal we want to reach. We must remember that what is shown in the lab and what is shown in the real world may be different. And we need to somewhat rein in our enthusiasm: just because it is “AI” it does not necessarily mean it is “better”. The future of AI lies mainly in decision support and efficiency assistance. We see it in other fields of life how efficient and helpful that can be, for grocery shopping, just to give an example. And only ten years ago we would have thought impossible to have such technology available in our everyday life.