

Toward knowledge-based liver surgery: holistic information processing for surgical decision support

K. März¹ · M. Hafezi² · T. Weller³ · A. Saffari² · M. Nolden¹ · N. Fard² · A. Majlesara² · S. Zelzer¹ · M. Maleshkova³ · M. Volovyk³ · N. Gharabaghi² · M. Wagner² · G. Emami² · S. Engelhardt¹ · A. Fetzer¹ · H. Kenngott² · N. Rezai² · A. Rettinger³ · R. Studer³ · A. Mehrabi² · L. Maier-Hein¹

Received: 20 February 2015 / Accepted: 20 March 2015 / Published online: 7 April 2015
© CARS 2015

Abstract

Purpose Malignant neoplasms of the liver are among the most frequent cancers worldwide. Given the diversity of options for liver cancer therapy, the choice of treatment depends on various parameters including patient condition, tumor size and location, liver function, and previous interventions. To address this issue, we present the first approach to treatment strategy planning based on holistic processing of patient-individual data, practical knowledge (i.e., case knowledge), and factual knowledge (e.g., clinical guidelines and studies).

Methods The contributions of this paper are as follows: (1) a formalized dynamic patient model that incorporates all the heterogeneous data acquired for a specific patient in the whole course of disease treatment; (2) a concept for formalizing factual knowledge; and (3) a technical infrastructure that enables storing, accessing, and processing of heterogeneous data to support clinical decision making.

Results Our patient model, which currently covers 602 patient-individual parameters, was successfully instantiated for 184 patients. It was sufficiently comprehensive to serve as the basis for the formalization of a total of 72 rules extracted from studies on patients with colorectal liver metastases or hepatocellular carcinoma. For a subset of 70 patients with

these diagnoses, the system derived an average of 37 ± 15 assertions per patient.

Conclusion The proposed concept paves the way for holistic treatment strategy planning by enabling joint storing and processing of heterogeneous data from various information sources.

Keywords Cognition · Decision support · Ontology · Liver cancer · Computer-assisted intervention · Treatment planning

Introduction

Primary liver cancer is one of the most common cancer types worldwide with an estimated half a million cases per year [1,3]. At the same time, attributed to its blood-filtering function, the liver is the second most common site for metastases, with some authors projecting a 40–50% chance of developing liver metastases from primary extrahepatic tumors [1]. The treatment for liver tumors is thus of enormous clinical importance. However, the choice of treatment is usually not obvious, as it depends on a wide range of factors. These include:

Patient-individual data represent all information that can be acquired for a patient for whom the treatment plan is prepared. This information can be (automatically) extracted from images (e.g., liver volume, tumor number, size, and location), laboratory reports (parameter values), or other sources of information (e.g., clinical reports, hospital databases, and genetic analyses). It can be related to the disease (e.g., diagnosis and recurrence), the liver anatomy and function (e.g., parenchyma health and bilirubin levels), or general information (e.g., age and habits). The amount and types of data acquired at different hospitals may differ considerably.

✉ K. März
k.maerz@dkfz.de

L. Maier-Hein
l.maier-hein@dkfz.de

¹ Department of Medical and Biological Informatics,
German Cancer Research Center, Heidelberg, Germany

² Department of General, Visceral and Transplantation Surgery,
University of Heidelberg, Heidelberg, Germany

³ Institute of Applied Informatics and Formal Description
Methods, KIT, Karlsruhe, Germany

Factual knowledge has been written down in quotable sources (e.g., clinical guidelines, studies, and educational books). This allows the physician to make predictions about the morbidity and mortality of the disease and the possible interventions on an objective basis. However, with an estimated amount of 3000 contributions per year regarding hepatocellular carcinoma (HCC) alone [9], it is impossible for a single surgeon to know all necessary data by heart. To address this issue, clinical guidelines make the amount of studies manageable by deducing general recommendations about recurring cases. It has been shown, however, that clinicians are often not familiar with written guidelines and thus are unable to observe them correctly during the actual care process [6]. This leads to an upper limit of guideline complexity. At the same time, highly condensed guidelines cannot appropriately map uncommon cases and complicated treatment plans, which is especially the case for multimodal treatments. As a result, guidelines typically give merely rough directions, taking into account only a fraction of patient-individual parameters (e.g., size and number of tumors), while detailed treatment decisions remain to the surgeon (e.g., whether a resection is performed in a one- or two-stage approach, if it is performed as an open surgical intervention or laparoscopically, etc.). Furthermore, once created, guidelines are static until updated and thus most of the time do not reflect the latest state of knowledge.

Practical knowledge results from experience. It comprises case knowledge that encompasses the ability to interpret patient-individual data, deduce a prognosis and implications for the treatment, as well as expert knowledge about treatment options and their respective strengths, weaknesses, and their practical application (e.g., one-stage or two-stage resection, and radiotherapy performed as a palliative, curative, adjuvant, or neoadjuvant treatment). Due to the limitations of clinical guidelines in providing treatment plans that are optimized for each individual patient (see previous paragraph), hospitals increasingly employ tumor boards consisting of a multidisciplinary team of experts (surgical oncology, medical oncology, radiation oncology, etc.) who derive (possibly multimodal) treatment plans, taking into account not only factual knowledge and patient data but in particular the practical knowledge contributed by the board members.

Overall, the process of treatment planning is thus highly complex, involves a variety of heterogeneous data, and depends crucially on the experience of the physicians in charge. On the other hand, the potential of computers for storing and processing large amounts of data has not yet been exploited. Hospitals have vast databases of liver cancer patients with known clinical outcome (*practical knowledge*), but these data are not yet systematically used for clinical decision making. Similarly, valuable *factual knowledge*, such as

new study results, is not automatically integrated into the process.

To date, the literature on computer-assisted liver tumor treatment has focused on surgical resection planning [7, 10, 12, 17] or intra-operative guidance for specific treatments [2, 13] by almost exclusively processing image data. Contributions regarding strategy planning usually focus on formalization and evaluation of guidelines [6, 16]. Currently, no concept allows to plan a treatment strategy multimodally and under consideration of all knowledge sources.

The long-term goal of our work, which was conducted within the *Transregional Collaborative Research Center 125 “Cognition-Guided Surgery”*, was to develop a system for knowledge-based navigated liver surgery that assists the physician in formulating a holistic treatment plan by making use of all available knowledge, including *patient-individual*, *practical*, and *factual knowledge* for the first time. This paper presents a concept for modeling, storing, and accessing all the data that are used throughout the process of patient diagnosis, treatment planning, therapy, and follow-up. In this context, we address the following research questions:

1. How can we model a patient?
2. How can we integrate factual knowledge, such as new clinical guidelines and studies, into the decision-making process?
3. How can we implement the infrastructure necessary for storing, accessing, and processing the data?

Materials and methods

As illustrated in Fig. 1, our concept is based on holistically processing of all the relevant data involved in the treatment planning process to support the physician with all the relevant information at the right time and thus to facilitate, optimize, and objectify clinical decision making. Implementing our vision requires a common semantic representation of patient data to handle *patient-individual data* as well as case data (*practical knowledge*) (“Modeling patient-individual data” section) representing *factual knowledge* on the basis of the patient model (“Formalizing factual knowledge” section) and providing the technical infrastructure (“Infrastructure” section).

Modeling patient-individual data

Core of the vision is a holistic patient model that comprises all the relevant information on a patient and can be stored, processed, and accessed throughout the clinical workflow. This is not only relevant for covering all the patient parameters referred to in studies or clinical guidelines but also for retrieving similar cases from the hospital database [*case-based reasoning* (CBR)] and for discovering new causal

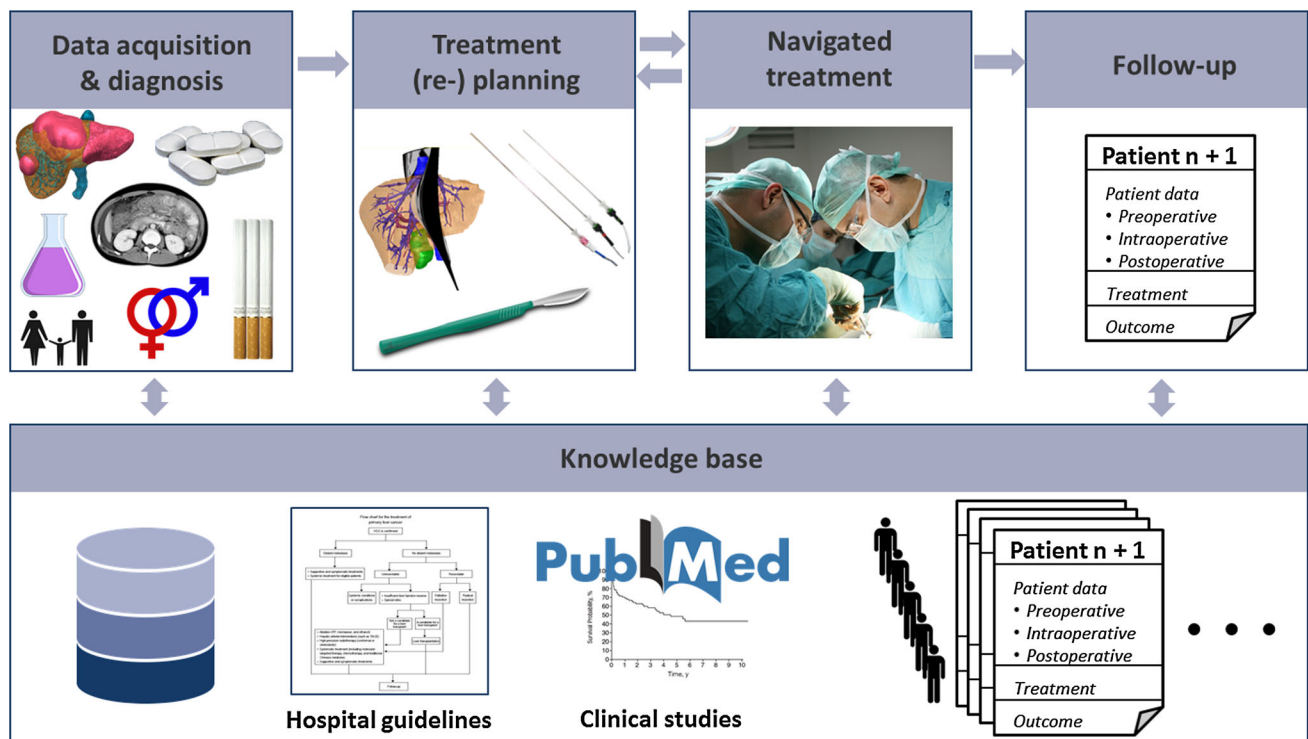


Fig. 1 Vision of knowledge-based liver surgery. The core component is the knowledge base, which dynamically acquires the patient data and all formalized knowledge. The treatment process comprises four stages. (1) Data acquisition and diagnosis: all the relevant patient data are collected and converted into a computer-interpretable format. Where possible, the system derives the data automatically from text or image sources. (2) Multimodal treatment planning: based on the acquired data,

the system processes the patient-individual data and the information in the knowledge base to derive a (possible multimodal) treatment plan. (3) Navigated treatment: the system assists the surgeon by providing adequate means of intra-operative navigation where necessary (which may be modified dynamically due to intra-operative finding). (4) Follow-up: in this phase, treatment outcome is documented, and the current case is fed back into the knowledge base

connections between parameters (*data mining*). We chose an iterative approach with the following steps to develop the model:

1. Collection of hospital-specific parameters: the patient model was initialized with all parameters that are collected for liver surgery patients in the regular clinical workflow at the University of Heidelberg.
2. Literature review: a team of four physicians performed a literature review, searching PubMed® for publications on liver tumor treatment published after 1980. A total of 213 relevant papers were selected, and factors related to diagnostics, treatment, or postoperative care in the context of liver tumor treatment were extracted.
3. Initial parameter structuring: the physicians were asked to organize the parameters into a form intuitive to them, while technicians supported the process. A provisional hierarchical model was created.
4. Refinement and prototyping: physicians started entering data from real-life cases into a form representing the preliminary model. In the process, usability, applicability, and completeness of the record were evaluated. In regular intervals, the data were reviewed and changes were made.

This process was repeated until changes became minor and largely consisted of introducing additional data rather than refining the model. At the same time, a computer-interpretable semantic representation for the data was developed. This representation was revised together with the model to achieve a structure that allowed to enter all types of data, represent it correctly, and make it available to the user in an intuitive manner.

5. Productive phase: a version was chosen to be used for large-scale data input for a retrospective case base (see “Patient model analysis” section). To ensure high data quality, we applied a three-stage supervision process, where data were double-checked by a colleague after each edit and a third time after a case was closed.

According to these steps, a model was derived to represent patient data. Most proposed systems are based on classical database structures [8, 14], which are unsuited for semantic modeling. Our aim was to create a model that enables and facilitates reasoning processes on our data. In this context, the resource description framework (RDF¹) has become an

¹ <http://www.w3.org/RDF/>.

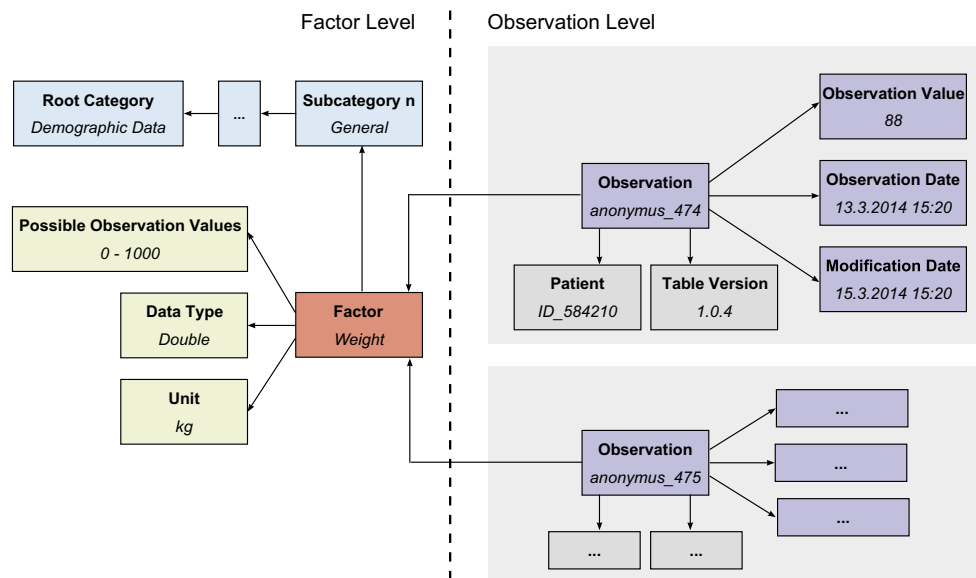


Fig. 2 Structure of the *factor* model. Central entities are the *factor* that models abstract knowledge about a *factor* and the *observation* that models specific characteristics of a certain patient. A *factor* has a specific data type and a unit that describes its nature. It also has a range of *possible observation values*, which might contain an interval for numbers or a list of strings. Each *factor* is also hierarchically organized, which allows for a tree-like resolution and a quick classification into a

primary context. The *observation* contains everything about a specific characteristic of a *factor* observed on a patient: the specific *observation*, the time it was observed, and the time it was entered or modified. It also contains information about the ontology version the *observation* was made under, and the patient this *observation* belongs to. More than one *observation* can exist for each *factor*, in which case the *observations* may form a time line

important tool because it enables a flexible formalization of information and provides the benefit that existing medical ontologies can be linked easily.

To date, we have identified a total of 980 relevant parameters that describe a patient in the context of liver cancer therapy. Each parameter is an atomic *observation* on that patient, but can be of arbitrary type (e.g., a CT image, the patient's gender, the fact that a type of surgery was performed, and a laboratory value). Often, some of these *observations* can be considered a subset of a more general clinical *Factor*. For example, the fact that a patient takes tetracycline and penicillin may be summarized under his prescription of antibiotics. In this way, *observations* were grouped into so-called *factors*. A *factor* is composed of several fields: a unique name, the data type, unit, and a range of *possible observation values*. See Fig. 2 for an example and for in-depth explanation. Our 980 parameters were grouped into 602 *factors*. As a detailed description would be beyond the scope of this paper, these factors will be made available on www.open-CAS.com and in a dedicated medical publication.

Additionally, *factors* are placed inside a category tree consisting of 23 root categories. Categories sort *factors* in a way that is comprehensible for the clinician. An example of how *factors*, categories, and additional information interact can be seen in Fig. 3. Additional semantic information can be stored for *factors*. This includes connections between *fac-*

tors, *observations*, and other abstract concepts (see Fig. 3). Overall, this design gives us a model with the following practical properties:

Heterogeneity: The abstraction of using a *factor* with required information in the form of a type characterization and a unit specification allows to insert arbitrary types of information and data. This way, the model is able to represent heterogeneous information, comprising, e.g., image data, past medical history, comorbidities, and treatment information. Additional semantic information can be conveyed by linking *factors* internally or to external sources (see Fig. 3).

Flexibility and extensibility: The data acquired for a patient are highly dynamic. For example, new imaging modalities and new laboratory parameters continue to be developed and acquired. The model must be able to cope with such changes over time. Adding a new *factor* requires only to add a name, data type, and unit to the model. Adding or creating new categories allows algorithms to automatically put the *factor* into a context. For example, data entry masks can display the new *factor* automatically to allow immediate input. This can be done with any ontology modeling tool or a simple Web-based tool provided with our model.

Multiplicity observations that have multiple instances or change over time (e.g., previous interventions or laboratory values) can be instantiated as often as necessary. Each instance has a time stamp for observation and modification,

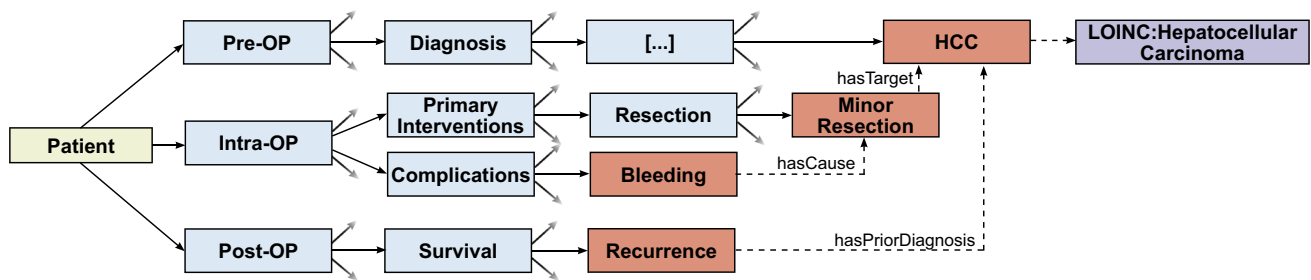


Fig. 3 Exemplary *factors* inside the category tree and additional information that links *factors* inside the model to outside concepts. Depicted here is an equivalence relation between our HCC *factor* and the LOINC² term of hepatocellular carcinoma

which allows to constrain analyses or reasoning to certain points in time.

Global applicability: Workflows and parameter usage differ between hospitals. Third parties can either use the presented structure or adapt the category trees and used parameters to their specific workflow and needs. This is helpful, e.g., if only a subset of *factors* can be used and if a certain laboratory test or procedure is not performed in a specific hospital.

Robustness: The system should be able to detect erroneous or conflicting input as early and often as possible. This is addressed by restricting data input to the defined unit, type, and possible *observation* values.

Usage of existing ontologies *factors* should be referenced to existing ontologies (e.g., LOINC² or SNOMED³). These ontologies act as a source for nomenclature, semantic hierarchy, and interoperability with other systems. *factors* can be linked to their counterparts in other ontologies to allow for interoperability (see Fig. 3).

Formalizing factual knowledge

Besides *practical knowledge*, which can be formalized using the patient model introduced in the previous paragraph (see also Fig. 2), another important source of information is *factual knowledge*, extracted from clinical guidelines or studies for example. This type of knowledge takes *factors* of a patient into account in order to deduce new *factors*. Examples of *factual knowledge* include a study that finds an increased risk of excessive bleeding in patients with a specific medication, an expert that classifies patients as being in “good general health” when they meet certain criteria or a widely accepted fact that a certain combination of antibodies found in the serum results in the diagnosis of an illness. We subsume statements made by all sources under the term *assertion*.

Based on a comprehensive analysis of *factual knowledge* from different sources (e.g., guidelines, studies, and litera-

ture), we decided to represent assertions in the form

$$(A?a) \wedge (B?b) \wedge \dots \wedge (X?x) \rightarrow Z = z \quad (1)$$

where capital letters represent a *factor* name of our patient model, the question mark implies an arbitrary comparison operator, and a lowercase letter implies a *possible observation value*. Subsequently, a focused formalization approach for specific cases was started to test the applicability of assertions on our data model. Since all derivations run on the proposed *factor* model, a reasoning system was developed that takes *factors* as input and produces *factors* as output. The result of this approach can be seen in Fig. 4. Most sources make assertions on patients with specific characteristics, which can usually be formalized using a single assertion. For these type of rules, we provide a tool that allows to generate formalizations semiautomatically by selecting *factors* and entering the conditions. However, more complex cases exist, such as in the case of decision trees. Our current approach is to normalize these trees by converting each edge between two nodes into an assertion of its own. The computer-readable interpretation of this assertion is then saved to a file of suitable format (see Fig. 4). The proposed model has the following properties:

Expressivity: By normalizing assertions to the given form, we were able to represent all *factual knowledge* we encountered during the study. The source class carries a minimal set of information that is required for all types of sources. Additional top classes can be chosen to define specific metadata (e.g., authors, journal, and evidence level for studies).

Significance: Not all sources are created equal. Trust put into results from a meta-analysis varies from that put in expert opinions or case reports. By using a central trust value, all sources have a common main variable to represent trust. Trust must then be calculated or assigned based on the source types. Studies, for example, are assigned a trust value based on their evidence level as proposed by the *National Guideline Clearinghouse*.⁴ In addition, each assertion can be assigned

² <http://loinc.org/>.

³ <http://ihtsdo.org/snomed-ct/>.

⁴ <http://www.guideline.gov/>, for example, see <http://hsl.lib.umn.edu/biomed/help/levels-evidence-and-grades-recommendations>.

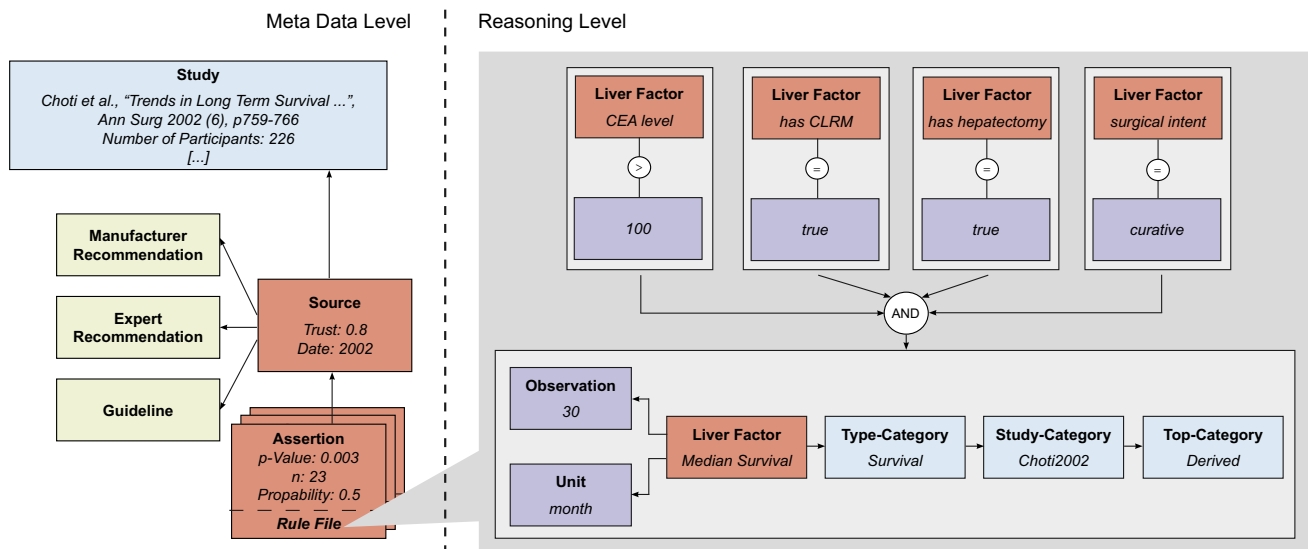


Fig. 4 Representation of *factual knowledge*. The *source* manages metadata regarding assertions. It contains common information to all sources: the date this source was published and the trust level that is associated with the source, which must be derived for each type of source individually. In addition, it has a body which defines its specific type in more detail. Each source has one or more assertions, which model a distinct piece of information from the source. The mode of deduction is the rule file, which contains a machine-interpretable query. The struc-

ture allows to query all rules eligible for a patient and retrieve their rule files, which are then executed. The resulting *factor* is then added into the category tree on the basis of his metadata. The assertion shown here states that a patient (1) with a CEA level above 100 (note that units are not saved in assertions), (2) with colorectal liver metastasis, (3) undergoing hepatectomy (4) with curative intent is expected to have a median survival of 30 months

a *p* value that represents its level of correlation inside the study.

Uncertainty: Most practical knowledge is uncertain, i.e., cannot be represented using a binary true/false value. As an example, a rule may predict that, under given circumstances, the probability that a given patient will develop liver failure will be 40 %. Each rule has a probability attached that models how probable the deduction is, if the assertion is considered true.

Multilevel reasoning: Deductions are always saved as instances of *factors*. As such, they can be used in further reasoning processes automatically; for example, based on the antibody presence, a Hepatitis B infection can be deduced, which in turn can lead to a negative effect on effective liver remnant volume.

Transparency: A deduced *factor* must always be identifiable as such. It must always be clear what led to the deduction of certain facts. Deduced *factors* are grouped in a parent category and form subcategories based on their source and type of statement. This way, *factors* can be grouped by their source automatically. Using the metadata stored by the source, it is always transparent how a piece of information was generated. Additionally, by using the explicit link to *factors*, all involved *factors* can be displayed as well.

Flexibility: Factual knowledge changes frequently and must be easily adaptable. This is of importance to be able to keep up

to date with new studies. Using our approach, new studies can be formalized at any time and are automatically integrated into the reasoning process. As soon as they are introduced, they can be included into the reasoning process.

Infrastructure

Application of our data modeling concepts for decision making in the clinical workflow requires a dedicated infrastructure to enable efficient data exchange and processing. Our concept for storing and accessing data is illustrated in Fig. 5 and is modeled along the following features:

Data storage: As mentioned above, the approach deals with highly heterogeneous data for modeling the patient. The storage system must not be limited in the data types that can be stored. However, certain common types like *Digital Imaging and Communications in Medicine* (DICOM) or RDF representations should be understood in particular by the system to allow for more structured access to it, e.g., to support multimodal and longitudinal relationships between data. To fulfill these requirements, the open-source imaging informatics software platform XNAT⁵ was chosen as a basis for our developments.

⁵ <http://www.xnat.org>.

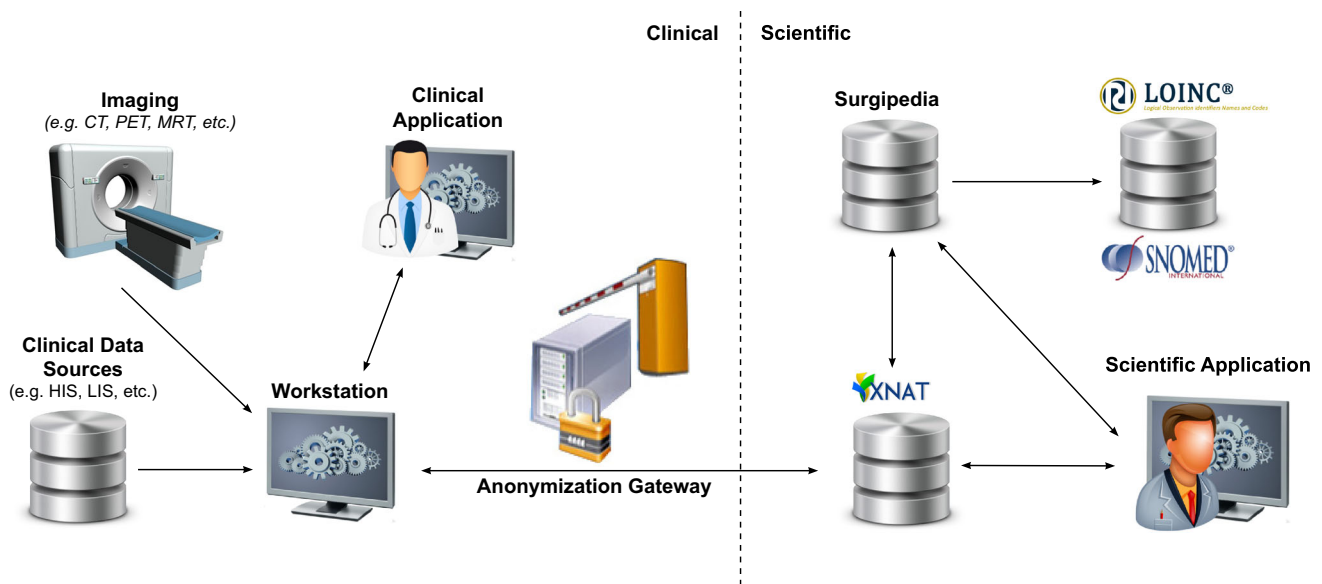


Fig. 5 Overview of the collaboration-wide infrastructure. Clinical personnel performs diagnostic procedures. The resulting image data and patient knowledge are stored anonymously on XNAT, where scientific projects can access the data

Patient data protection: To fulfill the legal requirements in an interdisciplinary project setting of clinical and non-clinical institutions, patient data must not be identifiable for non-clinical users of the infrastructure. Since data input comes from multiple sources at multiple time-points, a reproducible de-identification is necessary to merge the data of an individual case. To address this issue, we apply a DICOM de-identification service that generates reproducible subject IDs and feeds the data through XNAT’s DICOM gateway. Any other data can be uploaded by a *Representational State Transfer* (REST) API or interactively using a Web interface after a client-side de-identification. This ensures only de-identified data are kept in the system.

Data access: Data are accessed both interactively from users through various interfaces as well as by algorithms. To facilitate the deduction of new facts and the training of algorithms, flexible data access has to be ensured throughout the whole clinical and scientific workflow. By offering access to anonymized data through XNAT, external projects can access binary and semantic data for processing, algorithm training, and deduction of new facts (e.g., tumor progression prognoses or segmentations of liver and tumor). On the other hand, physicians have direct access to deanonymized data and can use the full data to plan operations for the patient.

Resource linking: An instance of the Semantic Media Wiki⁶ called *Surgipedia* provides central ontology design features and links to external ontologies. Additionally, we

developed a XNAT wrapper to semantically link data from XNAT and the *Surgipedia* in a manner transparent to the end user.

Experiments and results

For validation of our concept, we instantiated our patient model with real cases from the database of the Heidelberg University Hospital and formalized published studies using the proposed concept for handling factual knowledge. We then evaluated core usage statistics of our patient model (“Patient model analysis” section) as well as the capability of our concept to retrieve relevant patient-specific factual knowledge based on the proposed infrastructure for storing and accessing data (“Reasoning experiments” section).

Patient model analysis

To date, 184 patients are contained in the database with about 20–30 new patients being added each month. The database currently consists of a total of 84,475 *observations* (on average 459 *observations* per patient). Of the 602 *factors* identified in the development process (see “Modeling patient-individual data” section), 372 were covered by the patients instantiated so far, and 139 are used in more than 50% of the patients. An in-depth analysis of the distribution of *observations* and *factors* can be seen in Fig. 6. In a dedicated medical publication, the *factor* model and tools to create and export sets of *factors* will be made available to the public.

⁶ <https://semantic-mediawiki.org/>.

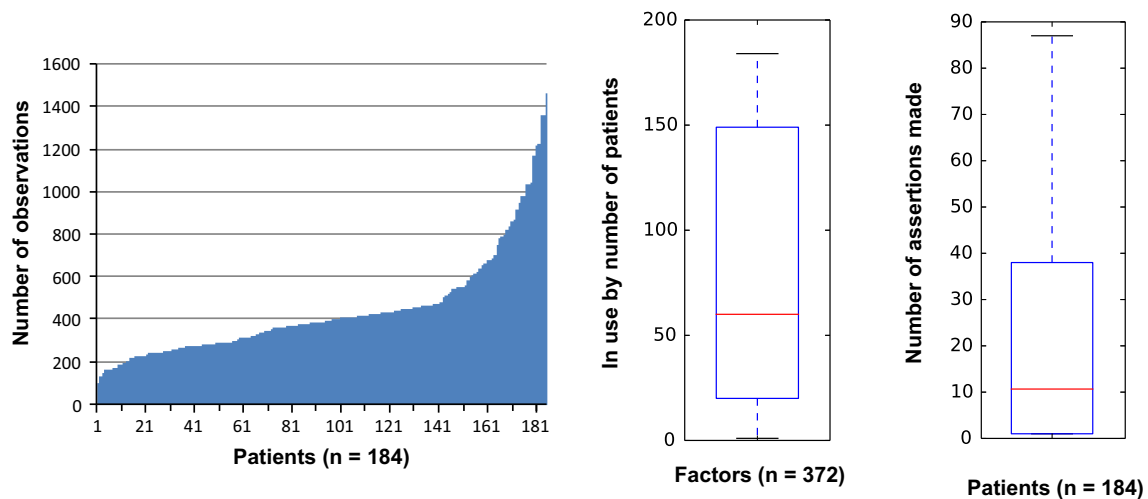


Fig. 6 Statistical analysis of our database. *Left*: histogram of the distribution of *observations* per patient. The center mainly consists of patients with common treatment scenarios. The peak to the right contains patients who had extensive laboratory testing procedures or more than one surgical intervention. *Center*: *factor usage* in our model.

Reasoning experiments

To show the applicability of our approach for modeling factual knowledge, in this initial feasibility study, we focused on mortality predictions for patients with either colorectal liver metastasis (CRLM) or HCC. These cases represent the most common primary and the most common secondary liver cancers [1] and could thus be expected to yield an acceptable coverage of our case base. The physicians in our project selected seven papers, which they regarded as relevant to the target patients and representative of available studies (i.e., differing in evidence level and evaluated assertions). We manually extracted the 72 assertions related to surgical outcome and survival.

Our patient model was sufficiently powerful to model all the content of the rules extracted. Their execution on the case base yielded a total of 70 matching patients who had been diagnosed with either HCC or CRLM. The system derived 38 ± 15 assertions per patient. In all 70 patients, we were able to generate data for a Kaplan–Meier curve to illustrate the predicted survival of the patient after the operation. Based on comorbidities like extrahepatic disease (EHD) and surgical outcome (e.g., R0-resection), additional curves could be generated that were either more specific for the patient (e.g., reduced survivability in case of extrahepatic diseases) or transported additional information (e.g., recurrence-free survival on basis of R0-resection). Additionally, several other probabilities were derived, e.g., chance for cure and perioperative mortality. A sample case for predictions made on a patient for CRLM resection can be seen in Fig. 7.

Shown are the number of patients that each *factor* is used in. *Right* number of derived assertions per patient. In both box plots, the boxes span the 25–75% quartile, and the *whiskers* mark minimum and maximum values

Discussion

The long-term goal of our work, which is conducted in the scope of the Collaborative Research Center TRR125 “Cognition-Guided Surgery”, was to develop a system that assists the physician in formulating a holistic treatment plan for liver surgery while making use of all available knowledge, including patient-individual, practical, and factual knowledge. To this end, we introduced a new concept for modeling, storing, and accessing all the data that are used throughout the process of patient diagnosis, treatment planning, therapy, and follow-up. These data include patient-individual data acquired from various different sources (e.g., imaging, laboratory, and demographic), practical knowledge (i.e., case knowledge), and factual knowledge such as clinical guidelines and studies. The following paragraphs discuss our contribution in the context of the state of the art.

Semantic patient model: Our model is flexible and extensible enough to adapt to the frequent changes in the clinical landscape. Placing the parameters in a category hierarchy allows us to organize parameters in a manner that is intuitive for physicians. While we currently have only one category hierarchy, *factors* can be placed in several distinct trees. This allows collaboration projects to reuse the *factors* and at the same time to create a *view* to our data, improving reusability. We are collaborating with groups working on radiation therapy, colorectal cancer resection, and heart surgery in order to adapt the model to their use cases and enable data integration. The database currently contains 980 parameters grouped in 602 *factors* that are deemed relevant by either the litera-

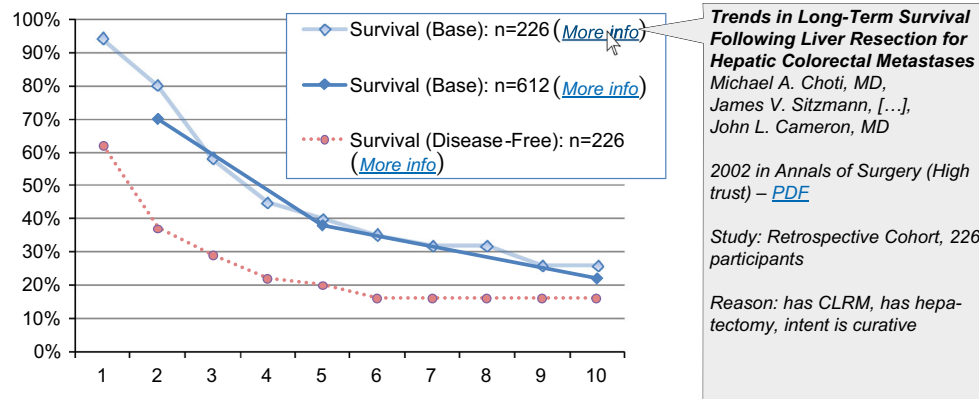


Fig. 7 Derived prognoses for a patient with a colorectal liver metastasis (CRLM) scheduled for hepatic segmentectomy according to our pilot study described in “Reasoning experiments” section. Two studies show the 10-year survival for patients under the same primary parameters (*base*), one of which has a lower case number but more data points [5], while the other has a higher case number and only three predictions [15]. The line transparency relates to the relative study importance (*currently the case numbers*). It can be seen that both studies predict similar prog-

noses for this patient. Additionally, a prognosis for disease-free survival was generated [5]. Additional facts that cannot be shown in graph form can also be displayed, e.g., perioperative mortality, chance for cure, and more specific median survival rates given specific outcomes of the operation. *Note* that contradictory prognoses are possible if different studies suggest different outcomes. In this case, both curves are shown in the plot, and the physician is able to assess their trust levels and their study parameters (see Fig. 7)

ture review or routine assessment in the University Hospital Heidelberg, Germany. To our knowledge, this is the largest collection of parameters for liver surgery that is currently available. Other contributions used a large set of parameters to establish a coefficient for patient similarity, but have not published the parameters themselves, so reuse is currently not possible (see., e.g., [8]). Currently, 57 % of the *factors* in the database are in use. One reason for this is that some of the *factors* refer to techniques that are not performed at our hospital, such as fluorescence-guided liver surgery. Another reason is the retrospective case selection: systematic input currently is focused on patients who underwent liver resection; accordingly, *factors* that are related to other therapies are not used in the model. A critical factor is data portability to ensure data validity throughout model updates. We track all changes to the model. In case *factors* change, equivalence relations are introduced to keep data comparable. Additionally, all data are time-stamped, which allows to constrain operations to certain time frames if historical versions should explicitly be ignored.

One core research question to be addressed in the context of incorporating practical knowledge in the decision-making process is how to define patient similarity based on our model. Previous approaches largely have used machine learning [4], image-based similarity [11], and known predictive factors [8]. These approaches are of limited applicability to our scenario because of the size of the problem domain (i.e., modeling of many therapeutical options) and data domain. Currently, we intend to exploit the category tree to calculate a semantic similarity between *factors*, which may allow us to combine current machine learning approaches with semantic constraints.

Factual knowledge: The second contribution is our concept for handling factual knowledge. The proposed framework allows to represent all relevant information and bring it into a machine-interpretable form. Previous approaches related to this challenge focused primarily on clinical guidelines [6, 16] and are not concerned with knowledge from other sources. Contributions from other domains, for example, match patient data with information about prescriptions to generate recommendations [4]. Standardized approaches for formalization are valuable, because they enable automatic import procedures. Accordingly, we intend to use results from related work if the data are available and applicable to us.

We have shown that our system is able to deduce relevant information about patients from studies. Kaplan–Meier plots have been generated for all patients matching our inclusion criteria. Additionally, more specific information was generated for patients depending on their *factors*, e.g., disease-free survival for patients with curative intent or reduced survival predictions for patients with extrahepatic disease.

Infrastructure: During data entry and processing, we actively used the proposed infrastructure. Use of the provided anonymization features allows us to share data throughout the collaboration. Additionally, the XNAT-Sever in combination with *Surgipedia* allows us to semantically link binary image data to the ontology. This is especially useful where research projects provide algorithms to derive additional data. An algorithm to segment the liver, for example, could be triggered by a CT image of the abdomen being uploaded to XNAT. The resulting segmentation is stored as well and annotated as belonging to the respective CT. Once a physician

marks the new factors as checked and valid, the results are available as annotated data to the treatment planning process and for algorithmic training.

Future work Future work should be devoted to the following aspects:

Large-scale data formalization: To investigate the clinical benefit of our method, we will focus on the implementation of clinical guidelines and improving study coverage. In this context, it would be interesting to establish an open collaboration base for this type of knowledge. By opening our collection of formalized data to the public and providing modeling tools for our framework, we take a first step in this direction. Furthermore, our long-term vision is to automatically formalize new studies published in PubMed via automatic text processing and understanding.

Integration of further information sources: An interesting aspect is the connection to other data sources. Establishing equivalence relationships between our terms and terms from other databases, such as the national database of liver transplants and genome maps, provides interoperability and completeness and offers opportunities for (semi)automatic data import. We constrained *factor* modeling to liver surgery in the first phase. With a growing number of additional information sources and import automation, applications in data mining become feasible. This would aid discovery of new interrelations in medical data.

Information aggregation: With an increasing number of studies, varying or even contradicting results will be more common. Aggregating studies that make different or even contradicting statements is a very difficult task. Results between two studies are rarely comparable, and while varying results are to be expected (e.g., source A predicting 32% chance of 5 year survival, source B predicts 41%), it is generally not possible to directly derive averaged values. Although possible from a modeling point of view, a statistically valid meta-analysis does not seem feasible in the current state but is an interesting outlook. However, with the trust level implemented, we have taken a first step to model study precedence and weighting. Still, more sophisticated approaches are being researched.

Practical knowledge: With a growing case base, extraction of knowledge from the database becomes more feasible. This includes methods working on images [11] and heterogeneous information [14].

Automatic data annotation: We intend to continue our systematic infrastructure integration into the clinical workflow. One focus is to automate and to integrate as much of the information acquisition as possible. At the same time, means

to automatically derive data as soon as new information is available are being developed.

In conclusion, we have taken first steps toward treatment strategy planning based on holistic processing of patient-individual data, practical knowledge, and factual knowledge. A validation study performed with 184 patients indicates that our concept for modeling, processing, and accessing data is suitable and paves the way for holistic treatment strategy planning by enabling the joint storing and processing of heterogeneous data from various information sources.

Acknowledgments This work was carried out with the support of the German Research Foundation (DFG) as part of project A02, I01, and S01, SFB/TRR 125 Cognition-Guided Surgery, with additional support from the projects A01 and B01. All of the authors state no conflict of interests. All studies have been approved and performed in accordance with ethical standards. Patient data were gathered and evaluated under informed consent only.

References

1. Ananthakrishnan A, Gogineni V, Saeian K (2006) Epidemiology of primary and secondary liver cancers. *Semin Intervent Radiol* 23(1):47–63
2. Beller S, Hunerbein M, Eulenstein S, Lange T, Schlag PM (2007) Feasibility of navigated resection of liver tumors using multiplanar visualization of intraoperative 3-dimensional ultrasound data. *Ann Surg* 246(2):288–294
3. Bosch FX, Ribes J, Daz M, Clries R (2004) Primary liver cancer: worldwide incidence and trends. *Gastroenterology* 127(5, Supplement 1):S5–S16
4. Chen RC, Huang YH, Bau CT, Chen SM (2012) A recommendation system based on domain ontology and SWRL for anti-diabetic drugs selection. *Expert Syst Appl* 39(4):3995–4006
5. Choti MA, Sitzmann JV, Tiburi MF, Sumetchotimetha W, Rangsin R, Schulick RD, Lillemoe KD, Yeo CJ, Cameron JL (2002) Trends in long-term survival following liver resection for hepatic colorectal metastases. *Ann Surg* 235(6):759–766
6. de Clercq P, Kaiser K, Hasman A (2008) Computer-interpretable guideline formalisms. *Stud Health Technol Inform* 139:22–43
7. Drechsler K, Oyarzun Laura C, Wesarg S (2012) Interventional planning of liver resections: an overview. In: *Engineering in Medicine and Biology Society (EMBC), 2012 annual international conference of the IEEE. IEEE*, pp 3744–3747
8. Dugas M, Schauer R, Volk A, Rau H (2002) Interactive decision support in hepatic surgery. *BMC Med Inform Decis Mak* 2(1):5
9. Gores GJ (2009) HEPATOLOGY: a home for hepatocellular cancer publications. *Hepatology* 50(1):1–2
10. Hansen C, Zidowitz S, Ritter F, Lange C, Oldhafer K, Hahn HK (2013) Risk maps for liver surgery. *Int J Comput Assist Radiol Surg* 8(3):419–428
11. Konukoglu E, Glocker B, Zikic D, Criminisi A (2013) Neighbourhood approximation using randomized forests. *Med Image Anal* 17(7):790–804
12. Mise Y, Tani K, Aoki T, Sakamoto Y, Hasegawa K, Sugawara Y, Kokudo N (2013) Virtual liver resection: computer-assisted operation planning using a three-dimensional liver representation. *J Hepatobiliary Pancreat Sci* 20(2):157–164
13. Peterhans M, vom Berg A, Dagon B, Inderbitzin D, Baur C, Candinas D, Weber S (2011) A navigation system for open liver surgery:

- design, workflow and first clinical applications. *Int J Med Robotics Comput Assist Surg* 7(1):7–16
14. Sun J, Wang F, Hu J, Edabollahi S (2012) Supervised patient similarity measure of heterogeneous patient records. *ACM SIGKDD Explor Newsl* 14(1):16–24
 15. Tomlinson JS, Jarnagin WR, DeMatteo RP, Fong Y, Kornprat P, Gonen M, Kemeny N, Brennan MF, Blumgart LH, D'Angelica M (2007) Actual 10-year survival after resection of colorectal liver metastases defines cure. *JCO* 25(29):4575–4580
 16. Wang HQ, Zhou TS, Tian LL, Qian YM, Li JS (2014) Creating hospital-specific customized clinical pathways by applying semantic reasoning to clinical data. *J Biomed Inform*
 17. Yang X, Lee W, Choi Y, You H (2012) Development of a user-centered virtual liver surgery planning system. *Proc Hum Fact Ergon Soc Annu Meet* 56(1):772–776