Decision Support for Data Segmentation (DS2):

Contextual Integrity Considerations

Martin French, Helen Nissenbaum, Mike Berry, Noam Arzt, and Carl A. Gunter

June 2, 2014

Acknowledgements

Thanks to The ILHIE Prototype Team (Mark Chudzinski, Daryl Chertcoff, Alexander Danel, Ivan Handler, Jean Kirupaharan, John Lekich, Antonios Michalos, Sireesha Perepu, Sam Nicolary, Maiko Minami, Vincent Bindschadler, Seeun Oh, Pete Sfiridis, Mario Felarca, Ting Wu, Thomas Redman, Bernie A'cs, David Stumpf, Tom Schultz, Ellick Chan, Jason Jacobs, Muhammad Naveed, and Sid Gupta). Thanks, also, to the New York University Information Law Institute Faculty and Research Fellows, and to the members of the New York University Privacy Research Group (thanks, especially to those who discussed this work with us: Solon Barocas, Catherine Dwyer, Travis Hall, Sophie Hood, Florencia Marotta-Wurgler, Ira Rubinstein, Luke Stark, Katherine Strandburg, Nathan Newman, Heather Patterson, Sasha Romanosky, and Malte Ziewitz).

This work was partially supported by HHS 90TR0003-01 (SHARPS). The views expressed are those of the authors only.

Contents

Abstract 3
1. Introduction 3
2. Data Segmentation for Privacy (DS4P) 6
2a. Background: The Concept of Data Segmentation 7
2a-i. Data Segmentation: A Mechanism for Increasing Patient Control Over Health Information Flows? 9
2a-ii. Data Segmentation: Some Key Challenges 9
3. The Contextual Integrity Framework: Assessing the Definition of Data Segmentation 10
3a. Two Approaches to Information Sensitivity: Context-Aware & Binary (a priori) 11
3a-i. The Context-Aware Approach 11
3a-ii. The Binary Approach 13
4. DS2: Building on DS4P Knowledge 14
4a. DS4P Use Case 15
4b. DS4P Pilots 16
5. DS2: Enabling Context-Aware Determinations of Sensitive Information 18
5a. A Brief Overview of Clinical Decision Support (CDS) Technology 18
5b. A Brief Overview of Clinical Document Architecture and Data Segmentation Challenges 19
5b-i. A Brief Introduction to the HL7 Clinical Document Architecture (CDA) & Continuity of Care Document (CCD) 19
5b-ii. Challenges Associated with Sequestering Information from a CCD 19
5c. Decision Support for Data Segmentation (DS2) 20
5c-i. The Core of the DS2 Architecture: Predicates and Reducers 20
5c-ii. Deterministic and Probabilistic Predicates 21
5d. Towards a Context-Aware Approach to Segmentation 22
5d-i. DS2 Deployed to Understand Context-Specific Concept-Relationships 23
5d-ii. DS2 Deployed to Understand Patient Expectations and for Patient Education 24
6. Conclusion 26

Abstract

Data segmentation is a concept that describes the process of sequestering elements in electronic records that are perceived as being undesirable to share in a particular context. In the face of increasingly ubiquitous flows of electronic health information, policy makers have called for the development of technologies that would enable data segmentation. Viewed as a means of giving patients more granular control over flows of their health information, data segmentation is thus a potential privacy-protecting strategy. Indeed, research into patient views of health information privacy and electronic health information flows suggests that patients want the capacity to segment their data. However, as data segmentation technologies are still nascent, there are important limitations associated with deploying data segmentation as a privacy-protection strategy. Chief among these limitations is that informed observers may be able to infer information that patients do not want to share, even after that information has been sequestered (we refer to this as the inferencing problem).

This White Paper proposes some novel approaches for addressing the inferencing problem. In conversation with the federally-sponsored "Data Segmentation For Privacy" (DS4P) initiative, it first discusses key challenges of data segmentation. It uses Helen Nissenbaum's Contextual Integrity framework to argue for the development of a more context-aware definition of data segmentation. It also discusses a research-oriented project demonstrating that certain DS4P tasks can be made more context-aware through the use of clinical decision support (CDS) technology. Dubbed **Decision Support for Data Segmentation** (DS2), this project advanced a unique use of CDS tools to (1) identify and sequester certain types of information in electronic medical records and to (2) mitigate potential risks of exchanging records from which data have been sequestered. It contributed a variety of tools that could be leveraged by existing DS4P pilot projects, including a new technical architecture and prototype, a suite of related open-source software tools, and test data for evaluating the application of various machine learning techniques to address the inferencing problem.

Our emphasis on contextual integrity considerations, and our engagement with the clinical inferencing problem, illustrate DS2's potential to help health information stewards and patients better understand the complex, context-aware behaviours of segmented records and corresponding data segmentation strategies.

1. Introduction

In the United States, electronic health information exchange (HIE) has assumed a central role in efforts to streamline and rationalize the organization of the health care system. Although the widespread digitization and exchange of health records is expected to have transformative effects at micro-, messo-, and macro-levels, it is recognized that such change "will create a new set of challenges for protecting the privacy and security of health information…".[1] For example, a 2006 report considering the development of a nationwide electronic health information network notes that many "individuals are concerned about the disclosure of their confidential personal health information because of possible embarrassment, emotional distress, and stigma. They are also concerned about more tangible harms, such as the inability to obtain employment, mortgages and

[1] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC). *Federal Health Information Technology Strategic Plan, 2011-2015* (Washington, D.C.: ONC, 2011): at 4.

other loans, or various forms of insurance".[2] While the disclosure of health records has always been accompanied by such potentialities, the ubiquitous exchange of electronic health information arguably creates conditions under which personal health information can be disseminated far more broadly than ever before. Harms arising from unwanted disclosures stand, therefore, to be potentially much greater. Thus, questions arise over how best to enable flows of personal health information while also addressing peoples' concerns about embarrassment, distress, stigma and discrimination, as well as building and preserving trust in systems of care. These questions are today confronted not just by HIEs, but also a host of other data stewards, including accountable care organizations (ACOs), hospitals, universities and others working with health information.

One response to such questions has been to advance the development of privacy-protective mechanisms in law, policy and technology. Thus, for example, the 2009 Health Information Technology for Economic and Clinical Health Act (HITECH)—part of the omnibus American Recovery and Reinvestment Act (ARRA)—included some significant expansions of extant privacy protections as specified by the Health Insurance Portability and Accountability Act (HIPAA).[3] HITECH also established, within the US Office of the National Coordinator for Health Information Technology (ONC), a Chief Privacy Officer to help articulate privacy policy, and to assure that privacy and security are addressed at every phase of technology development and implementation.[4] In addition, to help foster the development of privacy-protective technology, ONC has assumed responsibility for curating the standards and specifications that support interoperability amongst different electronic health record systems.[5] The federally orchestrated Standards and Interoperability (S&I) Framework provides a forum for HIE-related standards development.

These initiatives are underpinned, in large measure, by an ideal of individual autonomy. In this ideal, individual patients are given the means to grant or withhold the permission that ultimately underwrites flows of their personal information. The key mechanism underwriting personal information flows is the concept of consent. There are different ways of defining this concept. Christiansen and colleagues, for instance, have defined consent "as a process intended to determine the level of control of information disclosure and use in individuals."[6] For Goldstein and Rein,

---

[2] U.S. National Committee on Vital and Health Statistics (NCVHS), Subcommittee on Privacy and Confidentiality. *Privacy and Confidentiality in the Nationwide Health Information Network* [Internet]. Washington (DC): NCVHS; 2006 [cited 2014 March 19]. Available from: http://www.ncvhs.hhs.gov/060622lt.htm

[3] U.S. Government Printing Office (GPO). *American Recovery and Recovery Act* [Internet]. Washington (DC): GPO; 2009 [cited 2014 March 19]. Available from: http://www.gpo.gov/fdsys/pkg/BILLS-111hr1enr/pdf/BILLS-111hr1enr.pdf

[4] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC). *What is the Role of the Chief Privacy Office in the Office of the National Coordinator for Health Information Technology (ONC)?* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://www.healthit.gov/policy-researchers-implementers/faqs/what-role-chief-privacy-office-office-national-coordinator-heal

[5] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC). *Standards & Interoperability Framework* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://www.healthit.gov/policy-researchers-implementers/standards-interoperability-si-framework

[6] Christiansen, John, Apgar, Chris, and Melamed, Dennis (for the National Governors Association Center for Best Practices). *State and Federal Consent Laws Affecting Interstate Health Information Exchange* [Internet]. Washington, D.C.: National Governors Association Center for Best Practices, 2011 [cited 2014 March 19]. Available from: http://www.nga.org/files/live/sites/NGA/files/pdf/1103HIECONSENTLAWSREPORT.PDF

writing in the context of HIEs, consent can be understood as, simply, "patient permission to include personal health information…"[7] In the context of electronic health records, Kluge argues that, "in modern ethical thinking, informed, competent and voluntary consent is said to obtain when the patient has been given all the information that the objective reasonable person in the patient's position would want to know before making a choice [about who should have access to one's record]; further, that the patient must have understood the information at a subjective level, must also have understood the likely consequences of any choice that s/he might make, and must have made the choice in an authentic fashion".[8]

Clearly, it is no simple matter to answer the ethical question of whether or not a person's consent is informed, voluntary and authentic. Complicating factors still further is the fact that, in law and policy, different jurisdictions favour different consent mechanisms.[9] In the United States, for example, state laws differ from each other—and from federal regulations stipulated in HIPAA—with respect to how individuals must be involved in the disclosure of their health personal health information.[10] In addition, state and federal laws commonly prescribe that different classes of health information should be treated differently. For instance, information pertaining to treatment for substance abuse is typically regarded, in law, as de facto sensitive information. Taken together, these complexities are spurring the development new initiatives of processing, exchanging and protecting electronic health information.

In this paper we discuss one of the initiatives currently being developed under the auspices of the S&I Framework. Known as "Data Segmentation For Privacy" (DS4P), this initiative fosters the development of standards and technologies that are meant to enable actors to sequester certain data elements, which may be deemed undesirable to share, from an electronic medical record. While DS4P standards may help to enable health information exchange systems in which patients have more control over flows of their personal health information, we report on research that illustrates some of the limitations of segmentation as a privacy-protection strategy. Chief among these limitations is that informed observers may be able to infer information that patients do not want to share, even after that information has been sequestered (we refer to this as the inferencing problem). The inferencing problem may, in turn, prompt health organizations to withhold from exchange whole classes of information, even when they may have secured individual consent to share some of this information, under some circumstances.

Our analysis of data segmentation and the inferencing problem is informed by Helen Nissenbaum's Contextual Integrity (CI) framework.[11] Taking a values-in-design approach,[12] we used

---

[7] Goldstein, Melissa, and Rein, Alison. *Consumer Consent Options for Electronic Health Information Exchange: Policy Considerations and Analysis* [Internet]. Washington (DC): HHS ONC, 2010 [cited 2014 March 19]. Available from: http://www.healthit.gov/sites/default/files/choicemodelfinal032610.pdf

[8] Kluge, Eike-Henner. 2004. "Informed Consent and the Security of the Electronic Health Record (EHR): Some Policy Considerations," *International Journal of Medical Informatics* **73**(3): 229-234.

[9] Pritts, Joy, and Connor, Kathleen, *The Implementation of E-consent Mechanisms in Three Countries: Canada, England, and the Netherlands (The ability to mask or limit access to health data)* [Internet]. Rockville: Substance Abuse and Mental Health Services Administration, 2007 [cited 2014 March 19]. Available from: http://ihcrp.georgetown.edu/pdfs/prittse-consent.pdf

[10] Christiansen et al. *State and Federal Consent Laws Affecting Interstate Health Information Exchange*.

[11] Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy and the Integrity of Social Life*. Stanford: Stanford University Press; Nissenbaum, Helen. 2011. "A Contextual Approach to Privacy Online," *Daedalus* **140**(4): 32-48.

the CI framework to guide a research-oriented project, which demonstrated that certain DS4P tasks can be made more context-aware through the use of clinical decision support (CDS) technology. Dubbed **Decision Support for Data Segmentation** (DS2), this project advanced a unique use of CDS tools to (1) identify and sequester certain types of information in electronic medical records and to (2) mitigate potential risks of exchanging records from which data have been sequestered. It contributed a variety of tools that could be leveraged by existing DS4P pilot projects, including a new technical architecture and prototype, a suite of related open-source software tools, and test data for evaluating the application of various machine learning techniques to address the inferencing problem (the technical and architectural aspects of the DS2 project have been discussed in a companion white paper).[13]

Our emphasis on contextual integrity considerations, and our engagement with the clinical inferencing problem, illustrate DS2's potential to help health information stewards and patients better understand the complex, context-aware behaviours of segmented records and corresponding data segmentation strategies. For example, we suggest that DS2 technology could support research into patient privacy expectations and education initiatives that instruct patients and providers about the privacy-related implications of electronic health information exchange. Recognizing that, as a privacy-protective strategy, data segmentation is still in its infancy, we conclude by calling for the further development of communities of practice, organizations, and policy environments that clearly specify the appropriate flow of segmented health information, and which set out conditions under which it is appropriate to share segmented health records. Data segmentation alone, in other words, may be a necessary but is not a sufficient condition for the protection of personal health information in an era of ubiquitous electronic health information exchange. But, with appropriate technology and policy designed to foster new supportive assurances around DS2-enabled information exchange, privacy need not be put at risk by the emergent system of HIEs.

In the remaining five parts of this paper, our analysis unfolds as follows. In the next section—Section 2—we discuss the DS4P initiative, focusing especially on the working definition of data segmentation. In Section 3 we discuss the CI framework and offer a friendly critique of the data segmentation working definition. In Section 4 we discuss the DS4P Use Case Document and a DS4P Pilot demonstration project. In Section 5 we delve into an extended discussion of DS2, arguing that it enables a potentially more context-aware approach to data segmentation than anything else currently on offer. In Section 6, we conclude by calling for the further development of policy and technology that can build supportive assurances for the coming era of ubiquitous electronic health information exchange.

2. Data Segmentation for Privacy (DS4P)

As noted, the passage of the 2009 *Health Information Technology for Economic and Clinical Health Act* (HITECH) gave the ONC responsibility for the harmonization of health information technology

[12] Flanagan, Mary, Howe, Daniel, and Nissenbaum, Helen. 2008. "Embodying Values in Technology: Theory and Practice". In *Information Technology and Moral Philosophy*, Jeroen van den Hoven and John Weckert (eds). Cambridge: Cambridge University Press.

[13] Berry, Mike, Arzt, Noam, Chertcoff, Daryl, French, Martin and Gunter, Carl A. 2014. *Decision Support for Data Segmentation (DS2): Technical and Architectural Considerations*. Urbana-Champaign: SHARPS. Available from: https://sharps-ds2.atlassian.net/wiki/display/DS2/Publications.

(HIT) standards.[14] To fulfil its standards-related obligations, ONC established the S&I Framework in 2011. The S&I Framework is an open government process designed to orchestrate public- and private-sector input into the creation of harmonized HIT specifications.

The S&I Framework sponsors several initiatives, each designed to focus "on a single challenge with a set of value-creating goals and outcomes that will enhance efficiency, quality and effectiveness of the delivery of healthcare, through the development of content, technical specifications and reusable tools".[15] The S&I Framework maintains a wiki for coordinating amongst different stakeholders, and for archiving initiative materials.[16] Initiatives within the S&I Framework are vetted and prioritized by the S&I Framework Steering Team—made up of ONC leadership—and focus on challenges related, for example, to the electronic exchange of laboratory information, prescription drug monitoring, the creation of provider directories, and data segmentation for privacy" (DS4P).[17]

2a. Background: The Concept of Data Segmentation

Data segmentation may be defined as "the process of sequestering from capture, access or view certain elements that are perceived by a legal entity, institution, organization, or individual as being undesirable to share"—to qualify the working nature of this proposition, the DS4P wiki furthermore states: "This basic definition, however, does not account for the multiple permutations of segmentation in the health care context (i.e. granularity), nor does it adequately capture the varied considerations required for development of segmentation policy".[18]

---

[14] U.S. Congress. *Health Information Technology for Economic and Clinical Health (HITECH) Act*, Title XIII of Division A and Title IV of Division B of the American Recovery and Reinvestment Act of 2009 (ARRA), Pub. L. No. 111-5, Stat. 226 (Feb. 17, 2009), codified at 42 U.S.C §§300jj et seq.: §§17901. [Internet]. Washington (DC): US Congress; 2009 [cited 2014 March 19]. Available from: http://www.gpo.gov/fdsys/pkg/BILLS-111hr1enr/pdf/BILLS-111hr1enr.pdf

[15] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *Introduction and Overview* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://wiki.siframework.org/Introduction+and+Overview; see also: U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC). *Fact Sheets—Standards and Interoperability Framework* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://www.healthit.gov/sites/default/files/pdf/fact-sheets/standards-and-interoperability-framework.pdf; Halamka, John. "The ONC Interoperability Framework," *Life as a Healthcare CIO* [Internet]. 2010 March 30 [cited 2014 March 19]. Available from: http://geekdoctor.blogspot.com/2010/03/onc-interoperability-framework.html.

[16] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *Home* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://wiki.siframework.org/.

[17] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *S&I Initiative Overview, Phases and Outputs* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://wiki.siframework.org/S&I+Initiative+Overview,+Phases+and+Outputs.

[18] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *Data Segmentation for Privacy Charter and Members* [Internet]. Washington (DC): HHS ONC; n.d. [cited 2014 March 19]. Available from: http://wiki.siframework.org/Data+Segmentation+for+Privacy+Charter+and+Members. The source from which this definition has been adapted, is: Melissa and Alison Rein. *Consumer Consent Options For Electronic Health Information Exchange: Policy Considerations and Analysis* [Internet]. Washington (DC): HHS ONC; 2010. [cited 2014 March 19]. Available from: http://www.healthit.gov/sites/default/files/choicemodelfinal032610.pdf.

In theory, data segmentation could provide "a potential means of protecting specific elements of health information, both within an EHR [electronic health record] and in broader electronic exchange environments"—according to Jonathan Coleman, one of the DS4P coordinators, the "DS4P initiative addressed standards needed to protect those parts of a medical record deemed especially sensitive, or that may otherwise require additional privacy protection, while allowing other health information to flow more freely".[19] Accordingly, data segmentation and the DS4P initiative resonate with a broader vision in which patients will eventually be able to "control the flow of their data among institutions" and "choose what flows where for what purpose…".[20]

This broader vision, furthermore, stands as a potential technical solution to the challenges related to exchanging health information with organizations across jurisdictions. As Dimitropoulos and Rizk note, organizations wishing to share health information across jurisdictions "will have difficulty accommodating the current range of variation in policy requirements".[21] For instance, differences in the way that jurisdictions treat the process of obtaining consent could prove to be an impediment to cross-jurisdictional exchange. Organizations in states with opt-in approaches may be reluctant to share information with organizations in states that take an opt-out approach.[22] Similarly, variation in state law that specifies different treatment for different types of health information could be a barrier to widespread exchange. For example, in Illinois—where our project is based—state legislation "may restrict the ability of health care providers to make 'specially-protected' patient health information available to an HIE without having obtained a patient's prior 'opt-in' authorization. (The principle categories of 'specially-protected' patient health information in Illinois are: (i) mental health; (ii) substance abuse; (iii) HIV/AIDS; and (iv) genetic testing data)".[23]

A consequence of such variation is that organizations (and patients) may be reluctant to share entire documents with providers in jurisdictions that have different rules. As a result, extant HIEs may favour an "all-data-in or all-data-out out" suite of consent and sharing policies.[24] For this

---

[19] Coleman, Johnathan. "Segmenting Data Privacy: Cross-industry Initiative Aims to Piece Out Privacy Within the Health Record," *Journal of the American Health Informatics Management Association* 84(2): 34-38 [Internet]. Chicago: AHIMA. [cited 2014 March 19]. Available from: http://library.ahima.org/xpedio/groups/public/documents/ahima/bok1_050014.hcsp?dDocName=bok1_050014.

[20] Halamka, John. 2012. "Meaningful Consent," *Life as a Healthcare CIO* [Internet]. Blog [Cited 2014 March 19]. Available from: http://geekdoctor.blogspot.ca/2012/06/meaningful-consent.html.

[21] Dimitropoulos, Linda and Rizk, Stephanie. 2009. "A State-Based Approach to Privacy and Security for Interoperable Health Information Exchange," *Health Affairs* **28**(2): 428-434, p. 429.

[22] On variation in state consent models, see, for example: Dullabh, Prashila, Milstein, Julia-Adler, Nye, Christine, Moiduddin, Adil, Virost, Lindsay, Babalola, Elizabeth, Mahmud, Ayesha, and Jha, Ashish. *Evaluation of the State Health Information Exchange Cooperative Agreement Program: Early Findings from a Review of Twenty-Seven States* [Internet]. Chicago: NORC at the University of Chicago, 2012 [cited 2014 March 20]. Available from: http://www.healthit.gov/sites/default/files/pdf/state-health-info-exchange-coop-program-evaluation.pdf.

[23] Illinois, Data Security and Privacy Committee (DSPC), *Report of Preliminary Findings and Recommendations: State of Illinois Health Information Exchange Authority Data Security and Privacy Committee* [Internet]. Chicago: ILHIE (DSPC), 2012 [cited 2014 March 20]. Available from: http://sharps.org/wp-content/uploads/ILHIE_DSPC_Findings_091912_FINAL.pdf, p 7.

[24] Ibid. p. 9.

reason, a number of stakeholders have advocated the development of data segmentation technologies, which would enable the selective sharing of information in patient records.[25]

2a-i. Data Segmentation: A Mechanism for Increasing Patient Control Over Health Information Flows?

Ideally, data segmentation would allow patients to express complex consent preferences with respect to flows of their health information. They could consent, for example, to give their personal health information to state HIEs in order to facilitate flows of their information amongst organizations for the purpose of care. They could also consent to allow this information to be used for diverse secondary purposes, including research that could improve the quality and efficacy of care. In giving this consent, however, patients could opt to withhold certain elements of their electronic health record at certain points of time. They could opt, for example, to allow a nurse working at an immunization clinic to access portions of their electronic health record necessary for determining what vaccinations they require. At the same time, they could opt to prevent this nurse from reading portions of their record associated with their stay in an alcohol and drug abuse treatment facility.

2a-ii. Data Segmentation: Some Key Challenges

Realizing this broader vision of patient control over personal health information, however, will be challenging. Indeed, while data segmentation and the practice of sequestering information *could* yield greater patient control, it may not necessarily do so. Thus, as a privacy-protective strategy, the selective sharing of health information faces some key challenges. As Goldstein and Rein observe, for example, one key challenge is the determination of which information counts as "sensitive"—they argue that, although different types of data are recognized in law or policy as "sensitive," and therefore as requiring heightened protections or additional restrictions on disclosure, it may be difficult to determine *a priori* a) what information counts as an example of what might be legally recognized as sensitive, and, b) who (doctor? patient? hospital administrator?) should make this determination.[26] Organizations have tried different *a priori* approaches to determining information sensitivity. One approach, for example, is to construct a list of key words or coded concepts in EHRs that are automatically treated as sensitive. Another approach is to define types or classes of information as sensitive. This approach would treat all laboratory tests, for instance, as sensitive, or all medications as sensitive. A variation would be to class information according to source, and treat, for example, any information from federally funded substance abuse treatment facilities as sensitive.[27] Another issue is that patients may have a high-level opinion about what they consider to be sensitive but lack the knowledge to implement it in segmenting their own records, possibly because of the inadequacy of existing classifications.

Although list-based approaches go some way to enabling segmentation-based, privacy protective information-sharing strategies, they do not go far enough. They may be, in other words, necessary but not sufficient conditions for segmentation-based privacy protection. As we shall now argue, their key stumbling block is an inability to determine context-dependant sensitivity. Information sensitivity is a shifting target. To illustrate this point, we first describe in detail the framework of contextual integrity (CI). We then use the CI framework to consider how one might

[25] Illinois, Data Security and Privacy Committee (DSPC), Public Testimony [Internet]. Chicago: ILHIE (DSPC), 2012 [cited 2014 March 20]. Available from: http://www2.illinois.gov/gov/HIE/Pages/DataSecurityandPrivacy.aspx.
[26] Goldstein, Melissa and Alison Rein. *Data Segmentation in Electronic Health Information Exchange: Policy Considerations and Analysis*. Washington, DC: Office of the National Coordinator for Health IT, 2010.

[27] Ibid.

build upon and elaborate the DS4P working definition of data segmentation in order to more fully grasp the problem of determining information sensitivity.

3. The Contextual Integrity Framework: Assessing the Definition of Data Segmentation

Contextual integrity is an analytic framework for understanding and evaluating people's conceptions of a right to privacy.[6] Instead of claiming privacy as a right to control information about oneself, or to minimize access to it, the framework asserts that appropriate information flow is the touchstone for a right to privacy. It, further, posits context-specific informational norms as a model for appropriateness.

A key strength of the approach is that it provides a framework for determining when people are likely to experience a privacy violation. According to Nissenbaum, the underlying thesis of this framework is that "informational norms govern the flow of information about a subject from one party to another, taking account of the capacities (or roles) in which the parties act, the types of information, and the principles under which this information is transmitted among the parties"[28]. In other words, when information flows from point A to point B, there are multiple, independent variables that govern its movement. The subject of the information may have a say in what information flows where, to whom, with what frequency and under what circumstances, as may the sender and the recipient of the information whose access may be variously determined according to their roles in relation to the subject (e.g., doctor, nurse, family member). Beyond the immediate circle of these actors, however, lies an array of mediators. Administrative assistants, jurisdictional laws, organizational policies, computer-systems, and a host of other actors and factors help to constitute the norms that regulate information flows.

Taken together, the norms constituted in these structured social settings "prescribe and proscribe acceptable actions and practices".[29] In other words, they act as a set of implied rules for determining the (in)appropriateness of information flows. Indeed, according to the CI framework, the integrity of a context is preserved when norms are followed, and violated when norms are not. Furthermore, we would expect people to become upset—to experience a privacy violation—when context-relative informational norms are violated. Similarly, we would expect people to remain satisfied—to not experience a privacy violation—when context-relative norms are preserved.

Take, for example, a patient's expectation about flows of her personal health information. A patient may be satisfied—and even expect—her physician to share her health information with the surgeon she is about to see. But, she will likely be upset by—and would not expect—the sale of her personal health information by her physician to a marketing company. This is an example of a context-dependant expectation of privacy.[30] These expectations allow careful observers of informational norms to construct technologies and policies that are respectful of privacy and avoid heated public clashes over the privacy implications of novel interventions.

What does the CI framework tell observers about the working definition of data segmentation quoted above? In brief, the CI framework would ask the articulators of the working definition to go still further. What is necessary is *the delineation of the contexts in which data become sensitive, or undesirable to share*. As noted above, Goldstein and Rein observed that it may be difficult to determine *a priori* what information counts as sensitive information. This observation makes sense within the CI framework because information sensitivity is understood to vary with contextual

[28] Nissenbaum, *Privacy in Context*, p. 14.

[29] Ibid. p. 133.

[30] Nissenbaum, "A Contextual Approach to Privacy Online," p. 33.

parameters. For instance, if I have a new doctor, I may not want him/her to immediately know about my stigmatizing health information. However, I may consent to share this information with this person once we have established a good doctor-patient relationship. These kinds of nuances in understanding information sensitivity are not well captured by *a priori* classifications of information as either sensitive or not sensitive.

Given such examples, the CI framework stresses the need to build on the data segmentation working definition by discerning key contextual parameters. It challenges technological innovators to push beyond the patchwork of what is—or is not—permissible in law. In recognizing the potential of HIEs to radically alter health information flows, in recognizing the need of HIEs to establish relations of trust with patients, the CI framework emphasizes a nuanced approach to understanding information sensitivity. In particular, it will be important for data segmentation innovators to understand both established (codified) and emerging (contending) norms. These may be discerned by attending, at a minimum, to the following key parameters: 1) the key actors in the context (the information subjects, senders, and recipients); 2) the key attributes of the information flows (the types of information); and 3) the key transmission principles (the constraints that restrict information flows). Attending to these parameters will aid with the specification of norms operative in the context of health information exchange and with the context-aware determination of information sensitivity.

3a. Two Approaches to Information Sensitivity: Context-Aware & Binary (a priori)

As applied to electronic health information, articulations of the concept of information sensitivity have oscillated between proposals for context-aware determinations of what counts as sensitive versus proposals for binary, a priori determinations of what should be classified as sensitive.

3a-i. The Context-Aware Approach

Exemplifying the former approach, a founding document in the DS4P discourse on data segmentation reported on findings of an 18 month consultation process on privacy and confidentiality in the nationwide health information network (known then as the NHIN, and now as the NwHIN):

> Since the time of Hippocrates physicians have pledged to maintain the secrecy of information they learn about their patients, disclosing information only with the authorization of the patient or when necessary to protect an overriding public interest, such as public health. Comparable provisions are now contained in the codes of ethics of virtually all health professionals.
>
> As a practical matter, it is often essential for individuals to disclose sensitive, even potentially embarrassing, information to a health care provider to obtain appropriate care. Trust in professional ethics and established health privacy and confidentiality rules encourages individuals to share information they would not want publicly known. In addition, limits on disclosure are designed to protect individuals from tangible and intangible harms due to widespread availability of personal health information. Individual trust in the privacy and confidentiality of their personal health information also promotes public health, because individuals with potentially contagious or communicable diseases are not inhibited from seeking treatment.

> One of the major weaknesses of the current system of largely paper-based health records is its incomplete and fragmented nature. Ironically, this fragmentation has the unintended consequence of preventing disclosure of personal health information. Precisely because comprehensive health information is difficult to access, compile, use, and disclose, some health information privacy and confidentiality may be achieved by default.[31]

Here it is recognized that individuals may need to disclose information that they deem to be sensitive in order to receive appropriate care. And, it is noted that information systems contain transmission principles in their very architecture (in this case, paper-based systems place certain constraints on the disclosure of health information).

Furthermore, the report noted that health information is widely used for myriad purposes other than health care, and these uses are not always helpful for individuals. The report's justification of the need for context-aware determinations of information sensitivity bears quoting at length:

> Each year, as a condition of applying for employment, insurance, loans, and other programs, millions of individuals are compelled to sign authorizations permitting employers, insurers, banks, and others to access their personal health information for non-medical purposes. These authorizations are nominally voluntary; individuals are not required to sign them, but if they do not, they will not be considered for the particular job, insurance policy, loan, or benefit. […].
>
> An EHR system creates greater risks to confidentiality because the comprehensive disclosures might include much more information than is necessary to the particular decision at hand. At the same time, conversion to EHRs creates an unprecedented opportunity to protect confidentiality. At present, it may not be practicable to search a paper record system to disclose only a certain category of personal. Thus, personal [information] disclosed through compelled authorizations today is routinely overbroad, even where a narrower request is made. Conversion from paper records to EHRs could greatly enhance the confidentiality of personal health information and resolve the problem of excessive disclosures pursuant to authorizations. *Contextual access criteria* could be developed and integrated into the architecture of EHRs and the NHIN to permit disclosure of only the information needed by the user. For example, applying such technology, employers would only get information relevant to a particular job classification, and life insurers would only get information relevant to mortality risk. As a result, only personal [information] relevant to its intended use would be disclosed pursuant to an authorization.
>
> Developing the methodologies for these proposals will be complex and must involve collaboration by various stakeholders. The failure to incorporate contextual access criteria into the design of the NHIN, however, would have significant negative consequences, because this failure would impede the ability to limit unnecessary disclosures of irrelevant, sensitive personal [information] to third parties. Despite our certainty that contextual access

[31] U.S. National Committee on Vital and Health Statistics (NCVHS) *Letter to the Secretary - Recommendations Regarding Privacy and Confidentiality in the Nationwide Health Information Network* [Internet]. Washington (DC): NCVHS; 2006 [cited 2014 March 19]. Available from: http://www.ncvhs.hhs.gov/060622lt.htm.

> criteria are essential to protecting confidentiality in the NHIN, the NCHVS has been unable to identify any public or private research or pilot projects to develop this technology.[32]

The idea of developing contextual access criteria is well aligned with the CI framework. Yet the NCHVS did not go so far as to specify which parameters data segmentation innovators ought to consider in their development of policies and technologies that can assist in context-aware determinations of information sensitivity. Before discussing these parameters in greater detail, however, it will be worthwhile to consider the competing approach to determining information sensitivity.

3a-ii. The Binary Approach

In subsequent efforts to operationalize contextual access criteria, the notion that one should be able to "sequester" or withhold information from medical records became more tightly coupled with the notion that sequestered information was de facto "sensitive" information. For example, by 2009, Harry Reynolds (then) Chairman of the National Committee on Vital and Health Statistics recommended the following to Kathleen Sebelius, Secretary of the U.S. Department of Health and Human Services:

> PHR products should be designed to allow consumers to identify designated categories of sensitive health information. The consumer should then have the ability to control the use and disclosure of the information in these sensitive categories (including in emergency situations).[33]

This statement suggests a shift from contextually-determined sensitivity to categorically-determined sensitivity, a slippage that the CI framework indicates is to be avoided.

To better understand the potential limitations of the sensitive/non-sensitive binary, let us consider it in relation to a classic, binary distinction that has historically framed debates about privacy—the distinction between public and private information. As Nissenbaum notes, the public/private dichotomy is long-standing in discussions about privacy. Analysis of privacy discourse reveals at least three ways of conceptualizing this dichotomy. First, privacy is commonly conceptualized as a protective barrier between behaviour and regulation or control. Exemplifying this conceptualization is the distinction between the private lives of individuals and the actions of organizations, like government, to regulate the public (collective) lives of individuals. Second, in addition to being conceived of as a barrier that separates public from private, privacy is commonly conceptualized as a kind of sacred space, valuable in its own right for that which it allows to grow and develop. Third, applied to information, privacy is conceptualized as a form of access control. Taken together, these competing conceptualizations suggest that the private and the public are not independent variables; they are, rather, *inter-dependent* variables.[34]

In privacy scholarship, it has been commonplace to define information according to whether it is public or private (and therefore, in terms of whether it is not sensitive, or sensitive). As

---

[32] Ibid.

[33] U.S. National Committee on Vital and Health Statistics (NCVHS). *Letter to the Secretary – Protection of the Privacy and Security of Individual Health Information in Personal Health Records* [Internet]. Washington (DC): NCVHS; 2009 [cited 2014 March 19]. Available from: http://www.ncvhs.hhs.gov/090928lt.pdf.

[34] Nissenbaum, *Privacy in Context*, p. 90.

Nissenbaum observes, these positions assert that the proper place of privacy-related access protections is with respect to private information only.[35] Although the dichotomous distinction between public and private information has been rhetorically useful, its analytic traction has proven somewhat slippery for, as Nissenbaum argues, the line dividing public from private is neither static nor universal. Indeed, novel information technologies often prominently "reveal the inconstancy of boundaries and the fuzziness of definitions".[36] Moreover, the public/private dichotomy falls short in another important way. As Nissenbaum argues, because this view restricts privacy-related protections to pre-defined categories of private information, it neglects a range of situations—not typically deemed to be personal or private—in which there is a significant need for robust protections.[37] Accordingly, there is "a fundamental flaw in the notion that information can be divided into two categories, public and private, and that we need only worry about imposing constraints on the flow of private information".[38]

Instead of attending solely to the question of whether information is public or private, sensitive or non-sensitive, Nissenbaum argues that scholars, policy-makers and technology innovators must go further and examine the *context* in which information flows, and the context-dependant norms that both regulate information flows, and establish a base-line for expectations and judgments about the appropriateness of those flows. It is the teleology of a context—including its goals, purposes and ends—that inscribes with meaning the roles, activities, practices and norms that channel information flows. Adopting a CI framework, therefore, means going fundamentally beyond the dichotomous understandings of information given us by the public/private distinction, or by the sensitive/non-sensitive distinction. Instead of recognizing the importance of only two types of information, as these dichotomous understandings do, the CI framework "recognizes an infinite array of possibilities".[39]

To sum up, while the binary approach may provide a starting point for thinking about (and automating decisions about) information sensitivity, it is limited in some fundamental ways. Accordingly, there is a need to push policy and technology beyond the binary approach. Below we discuss research directed at the development of technology that can assist in the context-aware determination of information sensitivity.

4. DS2: Building on DS4P Knowledge

As noted, the DS4P initiative is articulating technical standards for the selective sharing of health information. It aims to foster the development of standards and technologies that are meant to enable actors to sequester certain data elements, which may be deemed undesirable to share, from an electronic medical record. As we shall now suggest, these standards—along with their operationalization in technical demonstration pilot projects—remain limited in their ability to execute context-aware determinations of information sensitivity. To illustrate this point we shall briefly discuss a DS4P Use Case Scenario and its operationalization in a DS4P pilot.

[35] Ibid. p. 96-98.

[36] Ibid. p. 101.

[37] Ibid. p. 114.

[38] Ibid. p. 120.

[39] Ibid. p. 143.

4a. DS4P Use Case

A goal of the DS4P *Use Case Document* is to help define the interoperability requirements for high priority data exchange within the S&I Framework. For instance, the DS4P initiative has as its goal the sharing of individually identifiable information to support 1) patient treatment and care coordination; 2) third-party payment; 3) health services research; 4) public health reporting; 5) population health; and 6) technology assessment and research. The requirements of such high-priority exchanges helped shaped the articulation of the user scenarios proposed in the DS4P *Use Case Document*. Moreover, the DS4P *Use Case Document* was written with a broad range of communities in mind—including patients, family members, providers, payers, vendors, standards organizations, public health agencies, and federal government bodies. It is intended to help work through the problem of implementing disclosure policies that originate from patients, from law, and from organizations.[40]

To illustrate, let us consider a scenario outlined in the *DS4P Use Case Document*. This scenario describes a push exchange (meaning that one care provider sends information—or pushes information—directly to another care provider).[41] In this scenario, a federally-funded Alcohol and Drug Abuse Treatment program (ADATP) sends information to a primary care provider (PCP). As described in the *Use Case Document*, the patient is notified that the ADATP is under legal obligation not to disclose patient information without the patient's consent. The patient indicates that he/she would like to disclose all of his/her treatment data to his/her PCP and completes an appropriate consent form authorizing the information exchange. The ADATP staff electronically captures the patient's information and consent directive, and the ADATP electronic system appropriately annotates the patient's record. The ADATP pushes the data, with an accompanying prohibition on re-disclosure notification, and the PCP's electronic system receives and incorporates the patient's annotated, protected data and prohibition on re-disclosure notification.

This scenario deals with the challenge of incorporating patient consent preferences into the electronic exchange of information. Other scenarios in the *DS4P Use Case Document* reflect more technically challenging problems for exchange, ranging from situations where patients want to change their information sharing preferences, to situations where those preferences are over-ridden in emergencies. Additionally, there are a number of issues that the DS4P *Use Case Document* acknowledges, but which the scenarios do not fully address. For example:

> Patients may provide conflicting consents that will be difficult to arbitrate electronically […]. There is little experience with the automated enforcement of prohibitions on re-disclosures […]. Data sent in response to queries may contain elements with different disclosure requirements, requiring that restrictions be communicated at a granular level.[42]

Moreover, the DS4P *Use Case Document* notes that there are risks associated with how data receivers interpret redacted data. Yet, as stated, these issues have been determined to lie beyond the scope of

---

[40] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC). *Use Case Development and Functional Requirements for Interoperability: Data Segmentation for Privacy* [Internet]. Washington (DC): HHS ONC, 2012, [cited 2014 March 19]. Available from: http://wiki.siframework.org/Data+Segmentation+for+Privacy+Use+Cases.

[41] "Push" exchanges are generally regarded as less technically complex than "pull" or query-based exchanges, which require an infrastructure for locating information on patients.

[42] HHS ONC, *Use Case Development*, pp. 50-51.

the DS4P initiative. Consequently, as we shall discuss, they remain ripe for future exploration in focused research studies that simulate for patients the implications of their segmentation choices. Before discussing these focused research studies in more detail, however, it will first be worthwhile to briefly discuss the operationalization of the DS4P use case scenarios in demonstration pilot projects, and also the DS2 approach to data segmentation.

4b. DS4P Pilots

To support work around the DS4P *Use Case Document*, the DS4P initiative has developed an *Implementation Guide*, and has approved a number of demonstration and pilot projects. The DS4P *Implementation Guide* is meant to "provide specific implementation guidance around the usage of standards and profiles" and "is designed to support developers and implementers who will be migrating to the recommended and evolving standards and technologies for data segmentation of healthcare information".[43]

The *DS4P Pilots* are meant to "exercise defined aspects of the Implementation Guide in a real-world setting," evaluating technology and standards while also providing "a test bed to evaluate the interaction of technology, implementation support, and operational infrastructure required to meet Data Segmentation Use Case objectives at the stakeholder or organization levels".[44] To illustrate how the pilots are operationalizing the technical standards associated with the DS4P initiative, it will be worthwhile to briefly discuss the first of the pilot demonstrations, which tackled the challenges associated with the push use case scenario.

Approved in June 2012, the VA-SAMHSA Pilot—a collaboration between the U.S. Department of Veterans Affairs, the U.S. Substance Abuse and Mental Health Services Administration, Mitre, Jericho Systems and HIPAAT—aims to build an open-sourced, extensible, access control service. This pilot intends to address all scenarios described in the DS4P *Use Case Document*. At the time of this writing, it had demonstrated a directed push exchange, sharing partial information, as outlined in scenario 1 of the DS4P *Use Case Document*, as well as a pull exchange, sharing all, as outlined in scenario 2 of the DS4P *Use Case Document*.[45]

Consider how this pilot demonstration addresses the requirements specified in the push use case scenario. It relies on meta-data tags to indicate the confidentiality, sensitivity and handling

---

[43] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *Data Segmentation Implementation Guidance, Version 1.0.3—Consensus Review* [Internet]. Washington (DC): ONC, 2012 [cited 2013 May 10]. Available from: http://wiki.siframework.org/Data+Segmentation+for+Privacy+IG+Consensus. The DS4P Implementation Guide is built upon several assumptions that, although outside of the scope of the present analysis, are nonetheless noteworthy. These stem from the choices made in the DS4P *Use Case Document* and include, for example, the assumption that patients "who are consumers of healthcare services are aware of their ability to complete Consent Directives and do offer such direction to the clinicians and organizations which they engage to provide them health care services" p. 14.

[44] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework. *Data Segmentation for Privacy RI and Pilots Sub-Workgroup* [Internet]. Washington (DC): ONC, 2012 [cited 2014 March 19]. Available from: http://wiki.siframework.org/Data+Segmentation+for+Privacy+RI+and+Pilots+Sub-Workgroup.

[45] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework, *DS4P VA-SAMHSA Pilot* [Internet]. Washington (DC): ONC, 2012 [cited 2014 March 19]. Available from: http://wiki.siframework.org/DS4P+VA-SAMHSA+Pilot.

instructions for exchanged information. As described during the demonstration (which is available online for review[46]), the process of exchange may be summarized as follows:

1. The sending organization must: A) verify that the receiving organization is authorized for the level of protected health information to be sent and determine restrictions that exist on use following delivery; B) provide document classification information as Access Control Decision Information and appropriately mark document sections with overall classifications; C) prepare a document set for delivery to the receiving organization, including handling instructions.

2. The receiving organization must: D) receive and process the document set; and E) verify that the receiving organization user is authorized to access the information according to organizational policy restrictions and information classification markings.

3. Both organizations must have access control services. The policies of the organizations and jurisdictions, and the patient's consent directives, become part of the access control information that is used to make decisions about what can be released. Resource access control information becomes information that is extracted for data tagging purposes.

4. The document is marked and put into an encrypted envelope; this is put in an outer envelope and sent. The envelopes are opened in turn; the access control information is put into the receiver's system. When an authenticated user attempts to get this data, access is governed by the access control information that has been sent with the document.

5. Trust between organizations is an absolute requirement for ensuring the privacy-protective features of this demonstration; however the establishment of supportive assurances and trusted relations amongst actors remains outside of the scope of the pilot, and of the use case scenarios.

Although the DS4P pilots are addressing important issues, it is noteworthy that some of the more challenging aspects of the scenarios set down in the DS4P *Use Case Document* have not yet been addressed. Many of the technical proposals for addressing the DS4P scenarios are dependent upon the appropriate annotation (for example, meta-data tagging) of exchanged records. Moreover, even assuming that exchanged records are appropriately tagged, the problem of segmenting data from records in such a way that sequestered information cannot be inferred by knowledgeable observers remains.

---

[46] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework, *The Standards and Interoperability (S&I) Framework in cooperation with Health Level Seven (HL7) Presents: Data Segmentation for Privacy* [Internet]. Washington (DC): ONC, 2012 [cited 2014 March 19]. Available from: http://wiki.siframework.org/file/view/HL7_DS4P_Presentation__2012_0909_FINAL.PDF/365435252/HL7_DS4P_Presentation__2012_0909_FINAL.PDF; see also: U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), Standards & Interoperability (S&I) Framework, *Data Segmentation for Privacy Initiative: All Hands Meeting* [Internet]. Washington (DC): ONC, 2012 [cited 2013 May 10]. Available from: http://vimeo.com/50334192

5. DS2: Enabling Context-Aware Determinations of Sensitive Information

To address some of the outstanding challenges faced by current approaches to data segmentation, a collaborative project involving researchers from the Strategic Healthcare IT Advanced Research Projects on Security (SHARPS), HLN Consulting, LLC, and the Illinois Office of Health Information Technology (OHIT), developed and prototyped a novel approach to determining information sensitivity. Grounded in the everyday operational challenges faced by the Illinois health information exchange (ILHIE), "Decision Support for Data Segmentation" (DS2) involves using clinical decision support (CDS) technology to help users understand the relations between data elements.

5a. A Brief Overview of Clinical Decision Support (CDS) Technology

CDS technology refers to computer-based systems designed to aid clinicians' decision making about individual patients. What distinguishes CDS technology from other forms of computer-aided analysis in medicine is the aim of providing support at the point of decision. In other words, while computer-aided analysis has a long been used to make retrospective sense of administrative and financial data, CDS technology is designed "to assist clinicians at the point of care".[47] Literature on CDS technology typically distinguishes between two different approaches. On the one hand are classic approaches that draw on a knowledge-base in order to provide relevant information to decision makers. On the other hand are approaches that employ machine learning and statistical pattern recognition techniques with the intent of overcoming some of the limitations of classic approaches.[48]

In the knowledge-base approach, decision support technologies typically involve two core components—a knowledge-base and a reasoning engine. Spooner describes these components in the following terms:

> The user supplies input appropriate to the system (i.e., terms from the system's controlled vocabulary to represent clinical data), and the system supplies output (e.g., a differential diagnosis). The reasoning engine applies formal or informal rules of logic to the input and often relies on additional facts encoded in the system's knowledge base. The knowledge base is the compilation of the relationships between all of the diseases in the system and their associated manifestations (e.g., signs, symptoms, laboratory and radiographic tests).[49]

As Spooner argues, a key challenge with such approaches is the maintenance of an up-to-date database, which must keep pace with changes in medical knowledge.

Accordingly, to address some of the challenges with the knowledge-based approach, researchers have developed machine learning and statistical pattern recognition techniques that are meant to help clinical decision support technologies learn from examples. These techniques allow systems "to learn over time," to change their behaviour "based on previous patterns".[50] However,

---

[47] Berner, Eta and La Lande, Tonya. 2007. "Overview of Clinical Decision Support Systems," in Eta Berner (Ed) *Clinical Decision Support Systems: Theory and Practice, Second Edition*. New York: Springer, p. 4.

[48] Ibid.

[49] Spooner, Andrew. 2007. "Mathematical Foundations of Decision Support Systems," in in Eta Berner (Ed) *Clinical Decision Support Systems: Theory and Practice, Second Edition*. New York: Springer, p. 35.

[50] Ibid., p. 40.

these techniques are not without their drawbacks, including, for instance, the requirement of substantial computing power.[51]

5b. A Brief Overview of Clinical Document Architecture and Data Segmentation Challenges

CDS technologies have, we argue, the potential to significantly aid the work of data segmentation. However, prior to the approach we are about to describe, no one had yet to our knowledge attempted to adapt CDS technologies to address the problem of data segmentation. Before we describe how CDS technologies can help with data segmentation, let us first explain why one would want to develop such an approach. Consider the following example.

A person living with HIV may not want to disclose her or his HIV serostatus to a chiropractor, deeming this potentially stigmatizing information unnecessary to a particular chiropractic treatment. Apprised of this information-sharing preference, the patient's primary care provider may seek to send certain information to the chiropractor while sequestering information related to the patient's HIV serostatus. In brief, the provider may wish to segment the patient's health record before sending it to the chiropractor. How could this process work using CDS technology? In order to describe DS2 as applied to this example, it will first be necessary to briefly describe the electronic documents that the provider would need to segment.

5b-i. A Brief Introduction to the HL7 Clinical Document Architecture (CDA) & Continuity of Care Document (CCD)

In the example described above, a patient's provider may seek to segment what is known as a Continuity of Care Document (CCD). CCDs are electronic documents that represent a clinical summary of a patient's encounter(s) with a provider. They are documents that are purpose-built for exchanging information amongst different providers. A typical health information exchange (HIE), for instance, might share CCDs amongst different providers as opposed to sharing a patient's entire set of medical records.

Contemporary CCDs[52] are built on an XML-based standard, and make up part of a broader structured document architecture known as the HL7 Clinical Document Architecture (CDA).[53] The CCD contains several standardized sections, including a section for allergies, and a section for medications, as well as a host of other structured sections relating, for instance, to a patient's encounters with different providers, their various medical problems, procedures they've had, diagnostic results, and so on. The CCD also contains header information, which includes information related to the subject of the document, for example, the patient's name.

5b-ii. Challenges Associated with Sequestering Information from a CCD

Taking our example introduced above, suppose a provider wanted to sequester information in a CCD that is related to a patient's HIV status. The provider would remove or redact the term "HIV" from the CCD's problem list. However, if this redacted CCD was subsequently sent to the chiropractor, the chiropractor could observe in the problem list that the patient has been managing

[51] Ibid., p. 41.

[52] Where the term "CCD" is used in this paper, it refers to the HL7 CCDA (HL7 Implementation Guide for CDA® Release 2: IHE Health Story Consolidation, Release 1.1 - US Realm, also known as the Consolidated CDA, or Consolidated Clinical Document Architecture) version of the CCD.

[53] Health Level Seven International (HL7). *HL7/ASTM Implementation Guide for CDA® R2-Continuity of Care Document (CCD®) Release 1* [Internet]. Ann Arbor: HL7, no date [cited 2014 March 19]. Available from: http://www.hl7.org/implement/standards/product_brief.cfm?product_id=6.

several opportunistic infections commonly associated with HIV. Moreover, the chiropractor could observe, in the medications section of the CCD, that the patient has been prescribed antiretroviral medication. The encounters section, furthermore, might disclose that the patient has been visiting clinics that are known to the chiropractor as clinics specializing in HIV treatment. Although none of this information in the CCD contains the term "HIV", it would nonetheless allow the chiropractor to infer that the patient in question is managing an HIV diagnosis.

What this example suggests is that the CCD must be comprehended in holistic terms. Because the different elements of the CCD are commonly related (albeit to varying degrees), a segmentation strategy based upon the redaction solely of terms in the record—in this case, the term "HIV"—would be insufficient for respecting the patient's information sharing preferences. Accordingly, what is required is a more context-aware approach to data segmentation, which takes into account, in this instance, the characteristics of the information-recipient, the characteristics of the information to be exchanged, and the transmission principles embedded in the CCD and regulating the exchange. DS2 has been designed to facilitate just such an approach. The overarching aim of DS2 is to help users understand when—assuming an informed observer—it will be possible to draw inferences about an element of information that a user wishes to sequester from sharing.

5c. Decision Support for Data Segmentation (DS2)

In the example we have just described, the segmentation strategy failed. It was not sufficient to respect the patient's information sharing preferences because it allowed the CCD recipient to infer that the patient was managing an HIV diagnosis. In fact, redaction of a condition and its related clinical facts often leaves residual facts—such as comorbidities and co-occurrences—that still reveal the condition.

To help address this kind of problem, the DS2 approach seeks to detect and segment not only the clinical facts that are directly related to a particular condition, but also the residual clinical facts—such as comorbidities and co-occurrences—that may still reveal the condition via inference. This is the inferencing problem we discussed in the introduction.

5c-i. The Core of the DS2 Architecture: Predicates and Reducers

DS2 is based upon a conceptual distinction between two different kinds of functions. The first kind of function is called a "predicate." Predicates are functions that are applied to a CCD, determining whether or not a specified property is present or absent in that CCD. The second kind of function is called a "reducer." Reducers recommend courses of action in the event that a predicate is triggered (e.g. in the event that a predicate determines the presence of a property).

As noted, the technical details and architecture of the DS2 approach have been described elsewhere.[54] However, the following description provides a brief illustration of how predicates and reducers work. Take, for instance, the problem of trying to sequester all information in a CCD related to an HIV diagnosis. Even focusing on only two sections of the CCD—the Problem List and the Medication List—it becomes quickly apparent that redacting HIV-related information is no simple task:

[54] Berry, Mike, Arzt, Noam, Chertcoff, Daryl, French, Martin and Gunter, Carl A. 2014. *Decision Support for Data Segmentation (DS2): Technical and Architectural Considerations*. Urbana-Champaign: SHARPS. Available from: https://sharps-ds2.atlassian.net/wiki/display/DS2/Publications.

Problem List:	**Medication List:**
• HIV infection • Candidiasis of lung • Bacterial infection, unspecified	• Combivir • Norvir • Procrit • Azithromycin • Fluconazole

Looking at the two lists above, the Predicate-Reducer system could be set up to identify and redact HIV infection. An HIV predicate would evaluate the Problem List and return "True" – an HIV condition is present. An HIV reducer would then remove the HIV diagnosis (it would redact "HIV infection").

But what about HIV medications (Combivir and Norvir)? An HIV predicate could be written such that it would evaluate the Medication List and return "True" because HIV medications are present. An HIV reducer could then remove the HIV medications ("Combivir" and "Norvir"). Yet, taking the two lists above, even after "HIV infection", "Combivir", and "Norvir" have all been removed, an informed observer may still be able to infer an HIV diagnosis. To address this possibility, an HIV predicate could evaluate the newly redacted CCD and return "True" again because of the co-occurrent clinical facts remaining in the record, which could allow observers to infer an HIV diagnosis.[55] Accordingly, an HIV reducer

> …might continue to remove co-occurrent clinical facts until the predicate returns "False." Ultimately, in addition to the HIV, Combivir, and Norvir, it might also remove the Candidiasis, Fluconazole, and Procrit – even though none of them are directly related to HIV. Alternatively, it might remove the bacterial infection, Azithromycin, and Procrit..[56]

Thus, what DS2 enables is an analysis of records for the presence of information that may be designated as sensitive, and a corresponding set of recommendations for how to redact a CCD to address the inferencing problem. Furthermore, DS2 enables an understanding of what we might think of as the "redactability" of records. That is, it enables users to determine the degree to which it will be possible to adequately sequester a piece of information in a given record, and thus to judge whether it is worth sending the redacted record in at all. Another way of saying this is that a predicate can be used to judge when the inferencing problem cannot (or should not) be solved by automated redaction.

5c-ii. Deterministic and Probabilistic Predicates

Given the above process, one can conceive of two different kinds of predicates: deterministic and probabilistic. Deterministic predicates leverage the knowledge-based approach to clinical decision support technology, whereas probabilistic predicates leverage pattern recognition and machine learning approaches.

[55] Ibid. p. 8, n13: "Candidiasis and Fluconazole may suggest an opportunistic infection and an immunocompromised patient. When this is combined with the bacterial infection and antibiotic, the suggestion may be stronger; and Procrit, a medication sometimes used to treat a side effect of the HIV medication, could add further support to the HIV inference. ".

[56] Ibid. p. 8.

With respect to deterministic predicates, it is possible to envision at least three classes. These classes vary in their complexity. Level 1 deterministic predicates are the least complex predicates. They are written to fire in the presence of obvious concepts. For example, if "HIV" or any clinical fact known to indicate or treat HIV is present, then fire. Level 2 deterministic predicates are predicates of moderate complexity. They are written to fire in the presence of correlated concepts. For example, if "HIV" is the target, then fire when comorbidities are present. Level 3 deterministic predicates are predicates of significant complexity. They are based on specific clinical rules. For example, if "HIV" is the target, then fire when two or more indirectly related concepts that suggest HIV are present; or, fire when a comorbidity is present and is consistent with an HIV-related laboratory result such as a CD4 count within a particular range.

The challenge with deterministic predicates is the difficulty in striking a balance between redacting too much and too little. If ***no*** co-occurring conditions are redacted, the door is left open to inferring a condition that is supposed to be hidden; if ***all*** possible co-occurring conditions are redacted, it is likely some conditions will have been redacted needlessly. A probabilistic approach can help with this challenge by applying real-world probabilities and machine learning techniques to develop predicates that learn to understand the density of ties between networked concepts. A Predicate-Reducer system based on such approaches can be optimized to redact the fewest number of clinical facts while still successfully preventing the inference of the targeted condition. For example, in the HIV example discussed above, the probability of inferring an HIV diagnosis after redacting Candidiasis, Fluconazole, and Procrit may be calculated to be lower than the probability of inferring HIV after redacting bacterial infection, Azithromycin, and Procrit. This may lead to a decision to redact the former set of facts as opposed to the latter set. But if the redaction of *only* Candidiasis and Fluconazole sufficiently lowers the probability of inference, it may be the preferred choice even if the inference probability after redaction is higher than the other choices.

5d. Towards a Context-Aware Approach to Segmentation

Reflecting on the analytic framework of contextual integrity, DS2 can potentially enable a nuanced, contextually-aware approach to segmentation of electronic health information. As noted, DS2 involves the use of a clinical decision support framework to help identify and redact information that could lead to unwanted inferences about a patient's condition. The CDS framework could deploy deterministic rules, as well as probabilistic rules, which consider the patient's clinical context, in addition to data about other patients and general clinical knowledge.

At its present stage of development, DS2 is focused primarily on data pertaining to what we might think of as the *clinical* domain. In other words, recommendations made by the DS2 prototype have mainly to do with clinical information (see Figure 1). According to the CI framework, this output must then be weighed against the ***rest*** of the context. Namely: 1) the key actors in the context (the information subjects, senders, and recipients); 2) the key attributes of the information flows (the types of information); and 3) the key transmission principles (the constraints that restrict information flows). Just as clinical decision support tools are designed to work in support of an ultimate arbiter (usually a human decision-maker), so too is DS2 set up to facilitate context-aware decisions about flows of patient information. In other words, the DS2 approach on the clinical side can be seen as analogous to what the contextual integrity framework calls for on the non-clinical side.

Figure 1

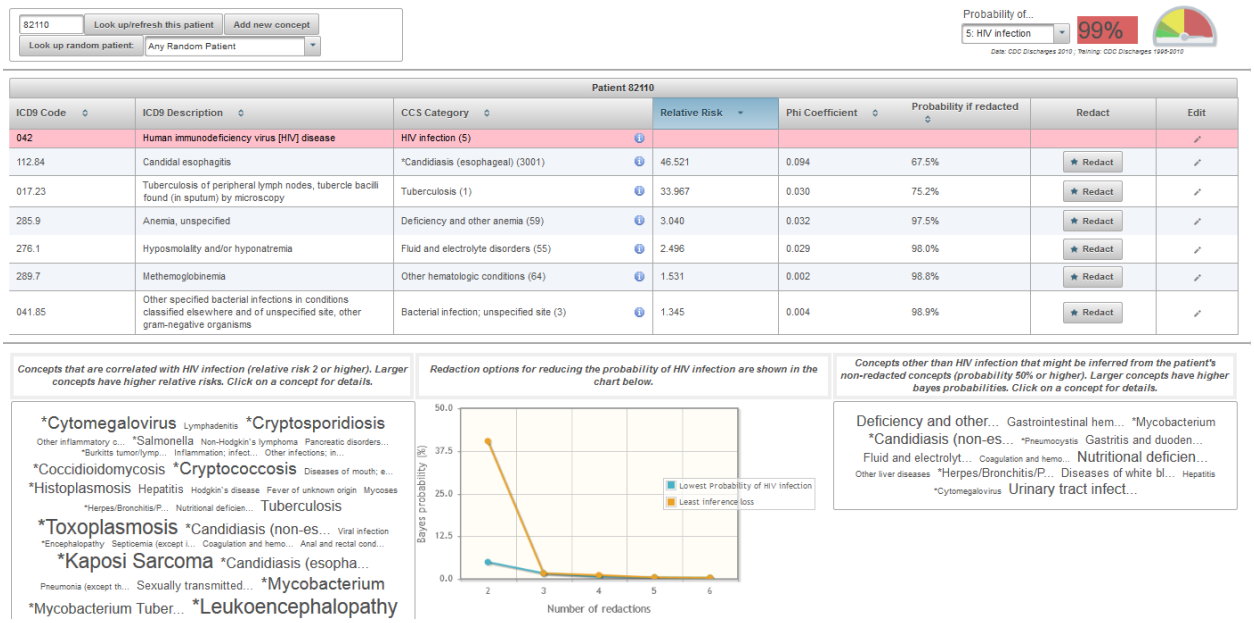

As DS2 technology continues to develop, it is possible to envision different ways that it could be deployed by HIEs and other data stewards. For example, aided by the decision support approach and the framework comprised by predicates and reducers, DS2 could be used to learn about the relationships between clinical concepts, particularly as these concepts transit through different inference-drawing contexts. Additionally, DS2 could be used to learn about 1) patient expectations about flows of their personal health information and 2) educate patients about the privacy-related implications of varying data segmentation strategies. We elaborate these different deployments below.

5d-i. DS2 Deployed to Understand Context-Specific Concept-Relationships

Recall that, according to the CI framework, a number of key factors and actors are at work in shaping the context-specific norms that govern information flows, and which determine the appropriateness of these flows. At a minimum, we can think about the following key parameters: 1) the key actors in the context (the information subjects, senders, and recipients); 2) the key attributes of the information flows (the types of information); and 3) the key transmission principles (the constraints that restrict information flows).

One potential deployment of DS2 could involve the articulation and testing of different predicate and reducer functions (rules). For example, it may be possible to represent different contextual parameters as distinct pools of knowledge, "dictionaries" of terms that group together related concepts, which collectively define higher-level, contextual domains.[57] We might think of these higher-level domains as providing a contextual basis for drawing inferences about concepts in a record.

[57] This approach is currently being developed by Dr. Ellick Chan and colleagues. See, for example, Chan, Ellick, Lam, Peifung, and Mitchell, John. 2013. "Understanding the Challenges with Medical Data Segmentation for Privacy," *USENIX Workshop on Health Information Technologies* [Internet]. Urbana-Champaign: SHARPS [cited 2014 March 31]. Available from: http://sharps.org/publications.

For instance, an information recipient could be conceptualized as having a specific set of knowledge that would make some kinds of inference-drawing extremely probable, and other kinds of inference-drawing extremely improbable. To illustrate, suppose we are sending a CCD to a haematologist. Because this recipient has specialized knowledge about haematology, it is reasonable to suppose that s/he could draw inferences that would not be possible for others (who do not possess the specific knowledge of a haematologist) to draw. In the case of a patient who has received an HIV diagnosis, therefore, one can envision differing strategies depending on whether an information recipient is a chiropractor or a haematologist.

Similarly, consider a consonant approach with respect to information attribute. A type of information, whether pertaining to a problem in the problem list or to an encounter in the encounters list, could be represented with respect to the varying density of ties it shares to the target concept. Information type, in other words, could be clustered with like concepts in ways that allow a person segmenting a record to understand how proximal or distal they are to target concepts. Likewise, transmission principles, whether pertaining to jurisdictional, institutional, or other types of restrictions on information flows, could be clustered according to the access-control rules they imply, and the inferences that these rules make possible.[58] With respect to the transmission principle example, for instance, a redacted CCD containing a specific prohibition on re-disclosure may allow an informed observer to infer that redacted information is "specially-protected information" under 42 CFR part 2, and therefore relates to a patient's treatment for substance abuse.

As we have argued in this paper, information sensitivity is more appropriately defined as a function of (changing) context than by a priori pre-determinations. The DS2 approach of distinguishing between predicate and reducer functions enables a nuanced consideration of the different contextual parameters that determine information sensitivity, in this case conceptualized as the probability that a recipient will be able to infer from a redacted document the information that a sender wishes to sequester. Predicates, in other words, enable a decomposition of context into (some of) its constituent parts. Reducers propose ways of treating CCDs given the different kinds of inferences made possible by different contexts (or different inferential configurations within contexts). DS2 provides a sandbox for articulating and understanding the implications of different predicate and reducer combinations.

5d-ii. DS2 Deployed to Understand Patient Expectations and for Patient Education

In a recent analysis of patients' desire for granular privacy control of their personal health information, focusing on which personal health information should be shared with whom, for what purpose, and whether these preferences vary based on the perceived sensitivity of health information, Caine and Hanania (2012) designed and undertook a card sort task with patients. Their study suggests that patients would prefer not to share all information (e.g., the current status quo) stored in electronic medical records with potential recipients.

Yet, in reflecting on what the study was able to accomplish, the authors observe a "significant limitation" of their study was that they "did not use personalized patient EMR data during the card-sort task".[59] The authors write:

---

[58] See, for example, Barth and colleagues' discussion of positive and negative norms of transmission: Barth, Adam, Data, Anupam, Mitchell, John, and Nissenbaum, Helen. 2006. "Privacy and Contextual Integrity: Framework and Applications," *IEEE Symposium on Security and Privacy* [Internet]. New York: IEEE [cited 2014 March 31]. Available from: http://www.andrew.cmu.edu/user/danupam/bdmn-oakland06.pdf.

[59] Caine, Kelly and Hanania, Rima. 2013. "Patients Want Granular Privacy Control Over Health Information in Electronic Medical Records," *Journal of the American Medical Informatics Association* **20**(1): 7-15

> Participants were given veridical examples of information that could exist in their EMR, but their data were not presented to them. A common finding from privacy literature in other domains is that de-contextualized sharing preferences often do not match actual sharing behaviour. It is unclear whether this mismatch is due to participants' inability to control sharing, lack of understanding about what is being shared and with whom, or whether *in situ* preferences differ from *a priori* preferences. Whatever the reason, the limitation to this study is clear: if participants had examined their own records they might have identified items they did not recall or know about previously, and this experience might have influenced their sharing preferences. In future studies, as well as in the design of privacy-enhanced EMRs it will be important for patients to see their own EMR data as they make privacy sharing decisions (8-9 –note omitted).

Caine and Hanania underscore the importance of being able to conduct experiments, with patients present, with their *own health information*. DS2 technology can facilitate exactly these kinds of experiments with patients.

HIE organizations and other data stewards (e.g. hospital records departments; bioinformatics units) are in potentially ideal positions to be able to deploy DS2 technology. In Illinois, the ILHIE has established expert consent management workgroups comprised of key stakeholders.[60] A key challenge in the workgroup discussions is to educate stakeholders on what kinds of segmentation strategies are technically possible, as well as which risks each strategy poses. DS2 technology could be deployed in such workgroup settings to help educate stakeholders on the implications of data segmentation. Workgroup participants would use test case health records that have been parsed with DS2 technology. This will indicate *in situ* how profoundly linked data elements in a record can be, as well as the degree of redaction necessary for truly redacting records such that even informed observers would not be able to infer information that patients might wish to withhold. It will also indicate the important role that communities of practice, organizations, and policy environments must play.

Additionally, the ONC has stressed that patient education and engagement "is an incredibly important part of meaningful consent"—according to the ONC's website on patient education and engagement:

> Education and engagement give patients information in an accessible and clear format so that they understand, for example:
>
> - what parts of health information could be accessed or shared
> - who could access health information
> - how health information is protected
> - why health information might be shared, and
> - the choices they have in terms of sharing or not sharing health information

[60] Illinois, Illinois Health Information Exchange (ILHIE), *Illinois Patient Consent Workgroups* [Internet]. Chicago: ILHIE, no date [cited 2014 March 19]. Available from: http://www2.illinois.gov/gov/HIE/Pages/PatientConsentManagementWorkshop.aspx

> Education and engagement are crucial to helping patients understand their consent options and *the impact of their consent choices* (emphasis ours).[61]

In our view, there is a potentially wide gulf between the work of providing patients with a menu of consent options, and the education required to equip patients with a working understanding of the impact of their consent choices. This is, perhaps, especially so with respect to patients who have complex engagements with systems of care, or with older patients who have less life experience with computer technology.

With respect especially to those patients who need the most help, we can envision the use of DS2 in specially designed counseling sessions that help them to make, in Caine and Hanania's words, "privacy sharing decisions".[39] In a document entitled *Points to Consider in Ethically Constructing Patient-Controlled Electronic Health Records,*[62] Meslin and colleagues explore the appropriateness of having special "information counselors" who would help patients understand the implications of their sharing choices. Although they adjudge this option ethically important, they ultimately rule it out on the basis of feasibility and cost to the system of care. We can envision utility, however, in model where patients who are interested in participating in HIEs, but who have complex health histories, special needs, or specific privacy concerns, could be given the option of making an appointment with their HIE to learn about technologically enabled privacy-protective strategies, like DS2.

In our view, this model could provide HIEs the opportunity to present patients with their own CCD, and also with what that CCD looks like after it has been redacted by DS2. This presentation would need, of course, to be embedded within a plain-language counseling session that would help patients to understand what information is contained in CCDs, who has the capacity in the HIE environment to access them, and under what circumstances. By engaging in such counseling, HIEs not only potentially educate patients about the impact of their consent choices, but also advance the process of building trust in an era of ubiquitous health information exchange.

6. Conclusion

In conversation with the DS4P initiative, this paper has argued that while data segmentation may be a useful privacy-protection strategy, it has some key limitations. One of the most fundamental of these, which we might think of as the clinical inferencing problem, is that informed observers may be able to infer information that patients do not want to share, even after that information has been sequestered. We have drawn on the framework of contextual integrity in order to analyze this limitation, and also to advance an argument in favour of a nuanced, context-aware approach to data segmentation. This approach, called **Decision Support for Data Segmentation** (DS2), uses clinical decision support technology to assess health records for instances of information that may allow informed observers to infer that which patients wish to keep private.

In conclusion, we wish to note that DS2 will take us only so far. Indeed, a key lesson emerging from the effort to develop DS2 has been the illustration that a great deal of information can be gleaned by making inferences based on remaining, un-sequestered information in a record.

---

[61] U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology (HHS ONC), *Patient Consent for eHIE: Patient Education and Engagement* [Internet]. Washington (DC): ONC, no date [cited 2014 March 19]. Available from: http://www.healthit.gov/providers-professionals/patient-consent-electronic-health-information-exchange/patient-education-and-engagement

[62] Meslin, Eric, Alpert, Sheri, Carroll, Aaron, Odell, Jere, and Schartz, Peter, *Points to Consider in Ethically Constructing Patient-Controlled Electronic Health Records* [Internet]. Indianapolis: Indiana University Center for Bioethics, 2012 [cited 2014 March 19], p 2. Available from: http://hdl.handle.net/1805/2936.

This has illustrated that data segmentation alone is not a sufficient condition for the protection of personal health information in an era of ubiquitous electronic health information exchange. But, with appropriate technology and policy designed to foster new supportive assurances around DS2-enabled information exchange, privacy need not be put at significant risk by the emergent system of HIEs.