

Cohesion and Coherence in the UDRP

T Lee and D Hunter

University of Pennsylvania, The Wharton School
3730 Walnut Street, Philadelphia, PA 19104

D Orr

Vanderbilt University Law School
131 21st Avenue South, Nashville, TN 37203

Abstract

Where the Internet was once said to dissolve borders, today it defines an entirely new space, one that transcends traditional geographic, political, and social boundaries. Yet this new frontier does not exist in a vacuum; its denizens are not devoid of a past. In this study, we report on early work exploring the impact of regional bias on transnational policymaking. The Uniform Domain Name Dispute Resolution Policy (UDRP) of the Internet Corporation for Assigned Names and Numbers (ICANN) is introduced as a test case. Prior work has established a significant disparity in the adjudication of the UDRP as a transnational policy. In this work, we explore the influences of nationality on transnational policymaking by testing the hypothesis of national bias in UDRP adjudication. We consider both the parties to an arbitration decision as well as the panellist(s) deciding the case. As a proxy, we also test the role of the panellist(s) legal tradition and training (i.e. common law versus civil law).

Introduction

Tip O'Neil, a former Speaker of the House of Representatives of the U.S. Congress is often cited for having observed that "All politics is local." The same can be said of dispute resolution. Both lawmaking and its concomitant adjudication have long been bound by local context. Legal jurisdictions are defined by the same geographic boundaries that circumscribe their respective city, state, or national borders.

The Internet's ability to disregard traditional, geographic boundaries may ultimately defy the tyranny of proximity. Yet even as scholars pursue a more deliberate investigation of the Internet's impact on local social and political norms, little attention has been directed at the associated issue of adjudication. This research aims to explore the impact of Internet technologies on the administration and evolution of justice by focusing on adjudication mechanisms.

This paper presents preliminary work on a broader project to investigate the premise that the Internet's global nature gives rise to the need for transnational policymaking. Specifically, given the establishment of a uniform policy across nation-state boundaries, this research first asks whether the attendant adjudication of the policy is similarly transcendent or whether it reflects selective elements of local tradition. Second, this research asks whether Internet technologies for sharing decisions on a global basis might reinforce or counteract any local influences on adjudication.

In this paper, we examine the Uniform Domain Name Dispute Resolution Policy (UDRP) of the Internet Corporation for Assigned Names and Numbers (ICANN) as one instance of transnational policymaking [7]. Based upon the existing and growing body of UDRP decisions, we report on the question of local influences on trans-national adjudication. Preliminary results suggest that nationality does not yield a bias where arbitrators disproportionately favor parties from their own country. Interestingly, legal tradition (whether arbitrators receive common law versus civil law training) may affect outcome.

We begin with some background on the UDRP and review prior work. Next, we briefly describe innovations in information extraction and verification developed to facilitate this study. The preliminary empirical analysis presented next is based upon our automated information processing and explores the role of nationality and legal training on outcome. We conclude with future work.

Background

In late 1999 ICANN introduced the UDRP to address the problem of “cybersquatting.” Cybersquatting occurs when an individual or institution registers a domain name which is the same as or similar to a (usually well-known) trademark and the registration is made in bad faith because it involves ransoming the name, or a similar tactic. The UDRP is an adjudication mechanism with no direct connection to any national court system. Instead, dispute resolution providers, such as the World Intellectual Property Organization (WIPO) and the National Arbitration Forum (NAF), to resolve cases. Each provider is an independent body with the authority to hear and decide the cases brought before it. For example, it is not uncommon to find a Swedish arbiter resolving a dispute between an American company and a Korean individual. Consequently, the UDRP represents a novel opportunity to study trans-national adjudication as it emerges from, and is influenced by, the borderless Internet.

The UDRP has proven remarkably successful on a number of levels. It is quick and cheap. In the last three years, more than 6000 cases have been filed, and more than 5000 decisions rendered. That efficiency is not without controversy, however. One possible cost is the potential for bias.

In prior work, researchers hand-coded cases into a relational database for statistical analysis [4, 14]. They then assessed fairness by counting case outcomes or the preponderance of specific panellists on particular cases. Depending upon whether the decision was rendered by a panel of one or three arbitrators and whether defendants file a formal response to the complaint or default on the claim, decisions favoured complainants over respondents in 60 to 80% of all cases. By implicitly assuming an even initial distribution of cases, the tremendous disparity in outcomes led researchers to question the fairness of the arbitration process [3, 13]. The ability to choose both arbitration forum as well as panel size (one versus three arbitrators), and the assignment of arbitrators to panels further reinforces questions of bias [2], though commentators are split on the seriousness of the charges [9].

Although existing research on UDRP bias in adjudication has relied upon hand-coded databases, a great deal of research on automatically extracting information from text into relations [5, 16] already exists. Some work has even focused on legal documentation [1, 15]. However, little effort has been directed at verifying the accuracy of such methods. Prior work on verification relied upon statistical techniques that could flag the possible presence of errors without identifying where the errors occurred [8, 10, 12].

Our work makes no assumptions about the initial distribution of cases. Instead, in this paper, we use novel, automated techniques to populate and verify a relational database of arbitration decisions and two hypotheses of bias: whether nationality and whether legal training affect outcomes. In particular, ours is the first work that we know of to incorporate panellist characteristics into the analysis and to probe for root causes. Moreover, we developed novel extraction and verification techniques that are unique in their use of semistructured constraints to both extract data and localize errors [6, 11].

Automated extraction and verification

Since its creation in 1999, the UDRP has produced thousands of decisions distributed across six different arbitration authorities located on three different continents. Though all decisions are published on the Web in HTML, there is not even a common index shared by multiple resolution providers let alone a single, comprehensive database of cases or decisions.

We implemented a wrapper-based system architecture to integrate data from multiple on-line arbitration authorities into a single, comprehensive database for subsequent statistical analysis [6, 11]. Although no cases are currently encoded in XML and every arbitration authority has its own conventions for structuring and presenting decisions, we have manually elicited a common, underlying Document Type Definition (DTD). Based upon that DTD, we define a set of

semistructured constraints that hold within and between different documents (e.g. inverse and inclusion constraints). Applying these constraints in a hierarchical manner [8], we use regular expressions to extract values into relations.

More significantly, we use the constraints to identify errors in the extraction. The constraints serve as internal and external (e.g. foreign key) consistency checks on the data. Unlike statistical methods of verification, which indicate whether a test set distribution fails to conform to a training set distribution, constraint-based checking identifies specific errors. Identifying specific errors is particularly significant in applications such as ours where we would otherwise calculate our aggregate statistics by hand. Greater detail is contained in the referenced papers.

Empirical analysis

Based upon the database constructed using automated tools developed for that purpose, we can then begin to test for bias. The nationality hypothesis is that UDRP panellists are more likely to side with complainants or respondents who are of the same nationality as they are. Because a single case may involve more than one domain and a single case may involve more than one panellist, we calculated our statistics based upon how individual panellists voted on individual disputed domains. It is possible that certain panellists were swayed in their vote, either for or against in reaction to their colleagues. However, barring these psychological factors, based on results of a sample of 2944 disputed domains where the nationalities of the panellist and both litigants were known, the nationality hypothesis appears to fail. Results are summarized in Table 1 (below), showing the relative rates at which complaints were denied and domains transferred by shared nationality. Nationality does not appear to exert either a sword effect in favor of complainants or a shield effect in favor of respondents. Domain transfer rates of panellists reviewing litigants from their home countries do not differ significantly from the sample mean of 85 percent. One caveat to these results, is that UDRP cases are dominated by U.S. panellists and litigants. Americans accounted for 66 % of complainants, 58 % of respondents, and 52 % of panellists in the sample. This resulted in 1501 transfers, more than half the sample, where the panellist, complainant, and respondent were all of the same nationality. Other nationalities may be more inclined to favor litigants from their home nations, but there are simply too few cases from other nations to draw statistically valid inferences about favoritism. We take up this issue below.

TABLE 1: Summary of Case Outcome by Shared Nationalities*

Shared Nationality	% Denied	% Transferred
All participants are of the same nationality (n=1745)	12.8	86
Complainant and panellist (n=1837)	13.0	85.7
Respondent and panellist (n=1798)	12.8	85.8
Complaint and panellist but not respondent (n=92)	17.4	79.3*
Respondent and panellist but not complainant (n=53)	15.1	79.2*
Sample mean (n=2944)	13.2	85.0

Values do not equal 100%. Other outcomes accounted for less than 1.5% of the sample and are excluded. * Differences are not statistically significant and can be attributed to small sample sizes.

In contrast to the idea of national bias, the sample provides some evidence that panellists are influenced by the legal tradition of the national origin. Panellists from civil law traditions transferred domains 9% more than panellists from common law traditions. This difference was highly significant ($p < 0.00$). Results are summarized in Table 2.

TABLE 2: Outcome by Panellist Legal Tradition

Legal Tradition	% Denied	% Transferred
Common Law (n=2219)	14.6	83.5

Civil Law (n=586)	6.7	92.2
-------------------	-----	------

Values do not equal 100%. Other outcomes and mixed legal traditions excluded.

As noted above, our initial results appear vulnerable to the heavy U.S. representation in the sample. Upon revisiting the issue, it appears that American panellists (87.1% transfer rate, n=1228) appear to be slightly more likely to transfer domains from respondents to complainants than their foreign counterparts (83.1%, n=1274). Although the effect size is relatively small, it is highly significant ($p < 0.00$). Were we to exclude Americans, as noted earlier, there are too few cases where the panellist and *only one party* are from the same foreign nation. However, the subsample of non-US panellists does include 327 domains at issue where *both parties* are from the same foreign nation as the panellist. The transfer rate within this sub-sample is 95.6 percent, substantially higher than the mean for foreign panellists (84.5 percent) and the sample mean of (85.0 percent). This effect, however, is likely the result of differences between civil (97.2 percent, n=281) and common law panellists (91.4%, n=32) within the subsample of foreign same-nationality cases. This then leads us to question the impact of U.S. skew on the influence of legal tradition. However, the effect of legal tradition is nearly identical in a subsample of foreign panellists (civil law=92.25, n=540, common law=84.5, n=579).

Future work

The results reported here describe the preliminary stages of a collaboration between scholars in information management and legal studies. The ultimate goal is three-fold: to develop novel tools for managing and manipulating semistructured data such as legal decisions, to better understand the evolution of trans-national policymaking and adjudication, and to develop insights on the relationships between the two. In particular, we would ultimately like to understand how information technologies can serve as a coordination mechanism, either mitigating or exacerbating the role of different local influences on trans-national decision-making. To that end, we have identified several immediate next steps.

In information management, we aim to continue our work on constraint-based extraction and verification. Specifically, we are currently exploring the combination of semi-structured constraints and classifiers for extraction and verification. Moreover, our current work relies upon a priori knowledge of constraints. We are also beginning to research the application of data mining and machine-learning techniques to learning semi-structured constraints, which can also represent DTDs.

For the empirical analysis, we will first expand the current database to better address issues of skew from the heavy U.S. population in the dataset. As the dataset grows, we also anticipate applying data mining techniques to search for interesting patterns within the decisions. In particular, we seek to investigate the expansion of particular legal concepts with respect to policy references in the UDRP as signs of trans-national legal evolution. The differentiation of arbitrators along nationalities and legal traditions promises to provide unique insight into this new space. The introduction of new top-level domains will provide interesting data for how existing local and UDRP precedents are interwoven.

Finally, not only do we hope to further automate our archive but also to make it publicly accessible. In so doing, we hope to be able to, through subjective surveys and outcomes analysis, test the impact of uniform access to a universal repository of decisions as a coordinating mechanism.

References

- [1] S. Bruninghaus and K. Ashley, "Toward Adding Knowledge to Learning Algorithms for Indexing Legal Cases," ICAIL-99, Oslo, Norway, 1999.
- [2] A. M. Froomkin, "ICANN's Uniform Dispute Resolution Policy-causes and (Partial) Cures," *Brooklyn*

- Law Review*, vol. 67, pp. 608, 2002.
- [3] M. Geist, "Fundamentally Fair.com? An update on bias allegations and the ICANN UDRP," accessed: 10 July 2002, aix1.uottawa.ca/~geist/fairupdate.pdf.
 - [4] M. Geist, "UDRPinfo.com," accessed: 9 September 2002, www.udrpinfo.com.
 - [5] J. Hammer, H. Garcia-Molina, J. Cho, R. Aranha, and A. Crespo, "Extracting Semistructured Information from the Web," presented at Workshop on the Management of Semistructured Data (SIGMOD-97), Tucson, AZ, 1997.
 - [6] D. Hunter, T. Lee, D. Ong, and Y. Yang, "Wrapper Verification Using Constraints for Semistructured Data," presented at WITS02, Barcelona, Spain, 2002.
 - [7] ICANN, "Uniform Domain Name Dispute Resolution Policy," 5 February 2002 www.icann.org/udrp.
 - [8] C. A. Knoblock, K. Lerman, S. Minton, and I. Muslea, "Accurately and Reliably Extracting Data from the Web: A Machine Learning Approach," *IEEE Data Engineering Bulletin*, vol. 23, pp. 33-41, 2001.
 - [9] A. Kur, "UDRP," accessed: www.intellecprop.mpg.de/Online-Publikationen/2002/UDRP-study-final-02.pdf.
 - [10] N. Kushmerick, "Wrapper Verification," *World Wide Web Journal*, vol. 3, pp. 79-94, 2000.
 - [11] T. Lee and Y. Yang, "Combining Information Extraction and Verification Using Semistructred Constraints: Position Paper," presented at International Conference on Database Theory (ICDT) Workshop on Data Quality in Cooperative Information Systems, 2003.
 - [12] K. Lerman and S. Minton, "Learning the Common Structure of Data," presented at AAAI-2000, Austin, TX, 2000.
 - [13] M. Mueller, "Success by Default: A New Profile of Domain Name Trademark Disputes Under ICANN's UDRP," accessed: 2002, dcc.syr.edu/markle.
 - [14] M. Mueller, "The UDRP Tracking Project," accessed: 9 September 2002, dcc.syr.edu/marklepage.htm.
 - [15] E. Schweighofer and D. Merkl, "A Learning Technique for Legal Document Analysis," presented at ICAIL-99 International Conference on Artificial Intelligence and Law, Oslo, Norway, 1999.
 - [16] S. Soderland, "Learning to Extract Text-based Information from the World Wide Web," presented at Third International Conference on Knowledge Discovery and Data Mining (KDD-97), 1997.