Practical Robust Estimators for the Imprecise Dirichlet Model

Marcus Hutter

RSISE @ ANU and SML @ NICTA
Canberra, ACT, 0200, Australia
marcus@hutter1.net   www.hutter1.net*

January 2009

Abstract

Walley's Imprecise Dirichlet Model (IDM) for categorical i.i.d. data extends the classical Dirichlet model to a set of priors. It overcomes several fundamental problems which other approaches to uncertainty suffer from. Yet, to be useful in practice, one needs efficient ways for computing the imprecise=robust sets or intervals. The main objective of this work is to derive exact, conservative, and approximate, robust and credible interval estimates under the IDM for a large class of statistical estimators, including the entropy and mutual information.

Contents

Keywords

Imprecise Dirichlet Model; exact, conservative, approximate, robust, credible interval estimates; entropy; mutual information.

*A shorter version appeared in the proceedings of the ISIPTA 2003 conference [Hut03].

1 Introduction

This work derives interval estimates under the Imprecise Dirichlet Model (IDM) [Wal96] for a large class of statistical estimators. In the IDM one considers an i.i.d. process with unknown chances[1] π_i for outcome $i\in\{1,...,d\}$. The prior uncertainty about[2] $\boldsymbol{\pi}=(\pi_1,...,\pi_d)$ is modeled by a set of Dirichlet priors[3] $\{p(\boldsymbol{\pi})\propto\prod_i\pi_i^{st_i-1}:\boldsymbol{t}\in\Delta\}$, where[4] $\Delta:=\{\boldsymbol{t}:t_i\geq 0\,\forall i,\sum_i t_i=1\}$, and s is a hyper-parameter, typically chosen between 1 and 2. Sets of probability distributions are often called Imprecise probabilities, hence the name IDM for this model. We avoid the term *imprecise* and use *robust* instead, or capitalize *Imprecise*. The IDM overcomes several fundamental problems which other approaches to uncertainty suffer from [Wal96]. For instance, the IDM satisfies the representation invariance principle and the symmetry principle, which are mutually exclusive in a pure Bayesian treatment with proper prior [Wal96].[5] The counts n_i for i form a minimal sufficient statistic of the data of size $n=\sum_i n_i$. Statistical estimators $F(\boldsymbol{n})$ usually also depend on the chosen prior: so a set of priors leads to a set of estimators $\{F_{\boldsymbol{t}}(\boldsymbol{n}):\boldsymbol{t}\in\Delta\}$. For instance, the expected chances $E_{\boldsymbol{t}}[\pi_i]=\frac{n_i+st_i}{n+s}=:u_i(\boldsymbol{t})$ lead to a robust interval estimate $[\frac{n_i}{n+s},\frac{n_i+s}{n+s}]\ni E_{\boldsymbol{t}}[\pi_i]$. Robust intervals for the variance $\mathrm{Var}_{\boldsymbol{t}}[\pi_i]$ [Wal96] and for the mean and variance of linear-combinations $\sum_i\alpha_i\pi_i$ have also been derived [Ber01]. Bayesian estimators (like expectations) depend on $\boldsymbol{t}$ and $\boldsymbol{n}$ only through $\boldsymbol{u}$ (and $n+s$ which we suppress), i.e. $F_{\boldsymbol{t}}(\boldsymbol{n})=F(\boldsymbol{u})$. The main objective of this work is to derive approximate, conservative, and exact intervals $[\min_{\boldsymbol{t}\in\Delta}F(\boldsymbol{u}),\max_{\boldsymbol{t}\in\Delta}F(\boldsymbol{u})]$ for general $F(\boldsymbol{u})$, and for the expected (also called predictive) entropy and the expected mutual information in particular. These results are key building blocks for applying the IDM. Walley suggests, for instance, to use $\min_{\boldsymbol{t}}P_{\boldsymbol{t}}[\mathcal{F}\geq c]\geq\alpha$ for inference problems and $\min_{\boldsymbol{t}}E_{\boldsymbol{t}}[\mathcal{F}]\geq c$ for decision problems [Wal96], where $\mathcal{F}$ is some function of $\boldsymbol{\pi}$. One application is the inference of robust tree-dependency structures [Zaf01, ZH05], in which edges are partially ordered based on Imprecise mutual information.

Section 2 gives a brief introduction to the IDM and describes our problem setup. In Section 3 we derive exact robust intervals for concave functions F, such as the entropy. Section 4 derives approximate robust intervals for arbitrary F. In Section 5 we show how bounds of elementary functions can be used to get bounds for composite function, especially for sums and products of functions. The results are used in

[1] Also called *objective* or *aleatory* probabilities.

[2] We denote vectors by $\boldsymbol{x}:=(x_1,...,x_d)$ for $\boldsymbol{x}\in\{\boldsymbol{n},\boldsymbol{t},\boldsymbol{u},\boldsymbol{\pi},...\}$, and i ranges from 1 to d unless otherwise stated. See also Appendix B.

[3] Also called *second order* or *subjective* or *belief* or *epistemic* probabilities.

[4] Strictly speaking, Δ should be the open simplex [Wal96], since $p(\boldsymbol{\pi})$ is improper for $\boldsymbol{t}$ on the boundary of Δ. For simplicity we assume that, if necessary, considered functions of $\boldsymbol{t}$ can and are continuously extended to the boundary of Δ, so that, for instance, minima and maxima exist. All considerations can straightforwardly, but cumbersomely, be rewritten in terms of an open simplex. Note that open/closed Δ result in open/closed robust intervals, the difference being numerically/practically irrelevant.

[5] But see [Hut07] for a proper Bayesian reconciliation of these principles.

Section 6 for deriving robust intervals for the mutual information. The issue of how to set up IDM models on product spaces is discussed in Section 7. Section 8 addresses the problem of how to combine Bayesian credible intervals with the robust intervals of the IDM. Conclusions are given in Section 9. Appendix A lists properties of the ψ function, which occurs in the expressions for the expected entropy and mutual information. Appendix B contains a table of used notation.

2 The Imprecise Dirichlet Model

This section provides a brief introduction to the IDM, introduces notation, and describes our generic problem setup of finding upper and lower statistical estimators. We first introduce the multinomial process and the Bayesian treatment with Dirichlet priors, and then the IDM extension to sets of such priors. See [Wal96] for a more thorough account and motivation.

Random i.i.d. processes. We consider discrete random variables $\imath \in \{1,...,d\}$ and an i.i.d. random process with outcome $i \in \{1,...,d\}$ having probability π_i. The chances $\boldsymbol{\pi}$ form a probability distribution, i.e. $\boldsymbol{\pi} \in \Delta := \{\boldsymbol{x} \in I\!R^d : x_i \geq 0\,\forall i,\ x_+ = 1\}$, where we have used the abbreviation $\boldsymbol{x} = (x_1,...,x_d)$ and $x_+ := \sum_{i=1}^d x_i$. The likelihood of a specific (ordered) data set $\boldsymbol{D} = (i_1,...,i_n)$ with n_i observations i and total sample size $n = n_+ = \sum_i n_i$ is $p(\boldsymbol{D}|\boldsymbol{\pi}) = \prod_i \pi_i^{n_i}$. The chances π_i are usually unknown and have to be estimated from the sample frequencies n_i. The maximum likelihood (frequency) estimate $\frac{n_i}{n}$ for π_i is one possible point estimate.

The Bayesian approach. A (precise) Bayesian models the initial uncertainty in $\boldsymbol{\pi}$ by a (second order) prior "belief" distribution $p(\boldsymbol{\pi})$ with domain $\boldsymbol{\pi} \in \Delta$. The Dirichlet priors $p(\boldsymbol{\pi}) \propto \prod_i \pi_i^{n'_i - 1}$, where n'_i comprises prior information, represent a large class of priors. The n'_i may be interpreted as (possibly fractional) virtual number of "observations". High prior belief in i can be modeled by large n'_i. It is convenient to write $n'_i = s \cdot t_i$ with $s := n'_+$, hence $\boldsymbol{t} \in \Delta$. Having no initial bias one should choose a prior in which all t_i are equal, i.e. $t_i = \frac{1}{d}\,\forall i$. Examples for s are 0 for Haldane's prior [Hal48], 1 for Perks' prior [Per47], $\frac{d}{2}$ for Jeffreys' prior [Jef46], and d for Bayes-Laplace's uniform prior [GCSR95]. From the prior and the data likelihood one can determine the posterior $p(\boldsymbol{\pi}|\boldsymbol{D}) = p(\boldsymbol{\pi}|\boldsymbol{n}) \propto \prod_i \pi_i^{n_i + st_i - 1}$.

The posterior $p(\boldsymbol{\pi}|\boldsymbol{D})$ summarizes all statistical information available in the data. In general, the posterior is a very complex object, so we are interested in summaries of this plethora of information. A possible summary is the expected value or mean $E_{\boldsymbol{t}}[\pi_i] = \frac{n_i + st_i}{n+s}$ which is often used for estimating π_i. The accuracy may be obtained from the covariance of $\boldsymbol{\pi}$.

Usually one is not only interested in an estimation of the whole vector $\boldsymbol{\pi}$, but also in an estimation of scalar functions $\mathcal{F} : \Delta \to I\!R$ of $\boldsymbol{\pi}$, such as the entropy $\mathcal{H}(\boldsymbol{\pi}) = -\sum_i \pi_i \log \pi_i$, where log denotes the natural logarithm. Since $\mathcal{F}$ is itself a random variable we could determine the posterior distribution $p(\mathcal{F}_0|\boldsymbol{n}) = \int_\Delta \delta(\mathcal{F}(\boldsymbol{\pi}) - \mathcal{F}_0) p(\boldsymbol{\pi}|\boldsymbol{n}) d\boldsymbol{\pi}$ of $\mathcal{F}$, where $\mathcal{F}_0 \in I\!R$ and $\delta()$ is the Dirac delta distribution. This may further be summarized by the posterior mean $E_{\boldsymbol{t}}[\mathcal{F}]=\int_\Delta \mathcal{F}(\boldsymbol{\pi})p(\boldsymbol{\pi}|\boldsymbol{n})d\boldsymbol{\pi}$ and possibly the posterior variance $\text{Var}_{\boldsymbol{t}}[\mathcal{F}]$. A simple but crude approximation for the mean can be obtained by exchanging E with $\mathcal{F}$ (exact only for linear functions): $E_{\boldsymbol{t}}[\mathcal{F}(\boldsymbol{\pi})]\approx\mathcal{F}(E_{\boldsymbol{t}}[\boldsymbol{\pi}])$. The approximation error is typically of the order $\frac{1}{n}$.

The Imprecise Dirichlet Model. There are several problems with this approach. First, the uniform choice $t_i=\frac{1}{d}$ depends on *how* events are grouped into d classes, which could be ambiguous. Secondly, it assumes exact prior knowledge of $p(\boldsymbol{\pi})$. The solution to the second problem is to model our ignorance by considering sets of priors $p(\boldsymbol{\pi})$, often called Imprecise probabilities. The specific *Imprecise Dirichlet Model* (IDM) [Wal96] considers the set of *all* $\boldsymbol{t}\in\Delta$, i.e. $\{p(\boldsymbol{\pi}|\boldsymbol{n}):\boldsymbol{t}\in\Delta\}$ which solves also the first problem. Walley suggests to fix the hyperparameter s somewhere in the interval [1,2]. A set of priors results in a set of posteriors, set of expected values, etc. For real-valued quantities like the expected entropy $E_{\boldsymbol{t}}[\mathcal{H}]$ the sets are typically intervals, which we call robust intervals

$$E_{\boldsymbol{t}}[\mathcal{F}] \in [\min_{\boldsymbol{t}\in\Delta} E_{\boldsymbol{t}}[\mathcal{F}]\, ,\, \max_{\boldsymbol{t}\in\Delta} E_{\boldsymbol{t}}[\mathcal{F}]].$$

Problem setup and notation. Consider any statistical estimator F. F is a function of the data $\boldsymbol{D}$ and the hyperparameters $\boldsymbol{t}$. We define the general correspondence

$$u_i^{\cdots} = \frac{n_i+st_i^{\cdots}}{n+s}, \quad \text{where } ^{\cdots} \text{ can be various superscripts or be empty.} \tag{1}$$

F can, hence, be rewritten as a function of $\boldsymbol{u}$ and $\boldsymbol{D}$. Since we regard $\boldsymbol{D}$ as fixed, we suppress this dependence and simply write $F=F(\boldsymbol{u})$. This is further motivated by the fact that all Bayesian estimators of functions $\mathcal{F}$ of $\boldsymbol{\pi}$ only depend on $\boldsymbol{u}$ and the sample size $n+s$. It is easy to see that this holds for the mean, i.e. $E_{\boldsymbol{t}}[\mathcal{F}]=F(\boldsymbol{u}\,;n+s)$, and similarly for the variance and all higher (central) moments. Most of this work is applicable to generic F, whatever it's origin – as an expectation of $\mathcal{F}$ or otherwise. The main focus of this work is to derive exact and approximate expressions for upper and lower F values

$$\overline{F} := \max_{\boldsymbol{t}\in\Delta} F(\boldsymbol{u}) \quad\text{and}\quad \underline{F} := \min_{\boldsymbol{t}\in\Delta} F(\boldsymbol{u}), \qquad \underline{\overline{F}} := [\underline{F},\overline{F}].$$

$\boldsymbol{t}\in\Delta \Leftrightarrow \boldsymbol{u}\in\Delta'$, where $\Delta':=\{\boldsymbol{u} : u_i\geq\frac{n_i}{n+s}\,\forall i,\, u_+=1\}$. We define $\boldsymbol{u}^{\overline{F}}$ as the $\boldsymbol{u}\in\Delta'$ which maximizes F, i.e. $\overline{F}=F(\boldsymbol{u}^{\overline{F}})$, and similarly $\boldsymbol{t}^{\overline{F}}$ through relation (1). If the maximum of F is assumed in a corner of Δ' we denote the index of the corner by $i^{\overline{F}}$, i.e. $t_i^{\overline{F}}=\delta_{ii^{\overline{F}}}$, where δ_{ij} is Kronecker's delta function, and similarly for $\boldsymbol{u}^{\underline{F}}$, $\boldsymbol{t}^{\underline{F}}$, $i^{\underline{F}}$.

3 Exact Robust Intervals for Concave Estimators

In this section we derive exact expressions for $\underline{\overline{F}}$ if $F:\Delta\to I\!R$ is of the form

$$F(\boldsymbol{u}) = \sum_{i=1}^{d} f(u_i) \quad\text{and concave}\quad f:[0,1]\to I\!R. \tag{2}$$

The expected entropy is such an example (discussed later). Convex f are treated similarly (or simply take $-f$).

The nature of the solution. The approach to a solution of this problem is motivated as follows: Due to symmetry and concavity of F, the global maximum is attained at the center $u_i=\frac{1}{d}$ of the probability simplex Δ if we allow $\boldsymbol{u}\in\Delta$, i.e. the more uniform $\boldsymbol{u}$ is, the larger $F(\boldsymbol{u})$. The nearer $\boldsymbol{u}$ is to a vertex of Δ, i.e. the more unbalanced $\boldsymbol{u}$ is, the smaller is $F(\boldsymbol{u})$. But the constraints $t_i\geq 0$ restrict $\boldsymbol{u}$ to the smaller simplex

$$\Delta' = \{\boldsymbol{u}\,:\, u_i\geq u_i^0\,\forall i,\; u_+=1\} \quad\text{with}\quad u_i^0 := \frac{n_i}{n+s},$$

which prevents setting $u_i^{\overline{F}}=\frac{1}{d}$ and $u_i^{\underline{F}}=\delta_{i1}$. Nevertheless, the basic idea of choosing $\boldsymbol{u}$ as uniform / as unbalanced as possible still works, as we will see.

Greedy $\boldsymbol{F(u)}$ minimization. Consider the following procedure for obtaining $\boldsymbol{u}^{\underline{F}}$. We start with $\boldsymbol{t}\equiv\boldsymbol{0}$ (outside the usual domain Δ of F, which can be extended to $[0,1]^d$ via (2)) and then gradually increase $\boldsymbol{t}$ in an axis-parallel way until $t_+=1$. With axis-parallel we mean that only one component of $\boldsymbol{t}$ is increased, which one possibly changes during the process. The total zigzag curve from $\boldsymbol{t}^{start}=\boldsymbol{0}$ to $\boldsymbol{t}^{end}$ has length $t_+^{end}=1$. Since all possible curves have the same (Manhattan) length 1, $F(\boldsymbol{u}^{end})$ is minimized for the curve which has (on average) smallest F-gradient along its path. A greedy strategy is to follow the direction i of currently smallest F-gradient $\frac{\partial F}{\partial t_i}=f'(u_i)\frac{s}{n+s}$. Since f' is monotone decreasing ($f''<0$), $\frac{\partial F}{\partial t_i}$ is smallest for largest u_i. At $\boldsymbol{t}^{start}=\boldsymbol{0}$, $u_i=\frac{n_i}{n+s}$ is largest for $i=i^{min}:=\text{argmax}_i n_i$. Once we start in direction i^{min}, $u_{i^{min}}$ increases even further whereas all other u_i ($i\neq i^{min}$) remain constant. So the moving direction is never changed and finally we reach a local minimum at $t_i^{end}=\delta_{ii^{min}}$. Below we show that this is a global minimum, i.e.

$$t_i^{\underline{F}} = \delta_{ii^{\underline{F}}} \quad\text{with}\quad i^{\underline{F}} := \arg\max_i n_i. \tag{3}$$

Greedy $\boldsymbol{F(u)}$ maximization. Similarly we maximize $F(\boldsymbol{u})$. Now we increase $\boldsymbol{t}$ in direction $i=i_1$ of maximal $\frac{\partial F}{\partial t_i}$, which is the direction of smallest u_i. Again, (only) u_{i_1} increases, but possibly reaches a value where it is no longer the smallest one. We stop if it becomes equal to the second smallest u_i, say $i=i_2$. We now have to increase u_{i_1} and u_{i_2} with same speed (or in an ε-zigzag fashion) until they become equal to u_{i_3}, etc. or until $u_+=1=t_+$ is reached. Assume the process stops with direction i_m and minimal u being $\tilde{u}$, i.e. finally $u_{i_k}=\tilde{u}$ for $k\leq m$ and $t_{i_k}=0$ for $k>m$. From the constraint $1=u_+=\sum_{k\leq m}u_{i_k}+\sum_{k>m}u_{i_k}=m\tilde{u}+\sum_{k>m}\frac{n_{i_k}}{n+s}$ we obtain $\tilde{u}=\frac{1}{m}[1-\sum_{k>m}\frac{n_{i_k}}{n+s}]=[s+\sum_{k\leq m}n_{i_k}]/[m(n+s)]$. One can show that $\tilde{u}$ as a function of m has one global minimum (no local ones) and that the final m is the one which minimizes $\tilde{u}$, i.e.

$$\tilde{u} = \min_{m\in\{1...d\}} \frac{s+\sum_{k\leq m} n_{i_k}}{m(n+s)}, \quad\text{where}\;\; n_{i_1}\leq n_{i_2}\leq ...\leq n_{i_d}, \quad u_i^{\overline{F}} = \max\{u_i^0,\tilde{u}\}. \tag{4}$$

If there is a unique minimal n_{i_1} with gap $\geq s$ to the 2nd smallest n_{i_2} (which is quite likely for not too small n and small s like 1 or 2), then $m=1$ and the maximum is attained at a corner of Δ (Δ').

Theorem 1 (Exact extrema for concave functions on simplices) *Assume $F:\Delta'\to I\!R$ is a concave function of the form $F(\boldsymbol{u})=\sum_{i=1}^d f(u_i)$. Then F attains the global maximum $\overline{F}$ at $\boldsymbol{u}^{\overline{F}}$ defined in (4) and the global minimum $\underline{F}$ at $\boldsymbol{u}^{\underline{F}}$ defined in (3).*

Proof. What remains to be shown is that the solutions obtained in the last paragraphs by greedy minimization/maximization of $F(\boldsymbol{u})$ are actually global minima/maxima. For this assume that $\boldsymbol{t}$ is a local minimum of $F(\boldsymbol{u})$. Let $j:=\mathrm{argmax}_i u_i$ (ties broken arbitrarily). Assume that there is a $k\neq j$ with non-zero t_k. Define $\boldsymbol{t}'$ as $t'_i=t_i$ for all $i\neq j,k$, and $t'_j=t_j+\varepsilon$, $t'_k=t_k-\varepsilon$, for some $0<\varepsilon\leq t_k$. From $u_k\leq u_j$ and the concavity of f we get[6]

$$\begin{aligned} F(\boldsymbol{u}')-F(\boldsymbol{u}) &= [f(u'_j)+f(u'_k)]-[f(u_j)+f(u_k)] \\ &= [f(u_j+\sigma\varepsilon)-f(u_j)]-[f(u_k)-f(u_k-\sigma\varepsilon)] \;<\; 0, \end{aligned}$$

where $\sigma:=\frac{s}{n+s}$. This contradicts the minimality assumption of $\boldsymbol{t}$. Hence, $t_i=0$ for all i except one (namely j, where it must be 1). (Local) minima are attained in the vertices of Δ. Obviously the global minimum is for $t_i^{\underline{F}}=\delta_{ii^{\underline{F}}}$ with $i^{\underline{F}}:=\mathrm{argmax}_i n_i$. This solution coincides with the greedy solution. Note that the global minimum may not be unique, but since we are only interested in the value of $F(\boldsymbol{u}^{\underline{F}})$ and not its argument this degeneracy is of no further significance.

Similarly for the maximum, assume that $\boldsymbol{t}$ is a (local) maximum of $F(\boldsymbol{u})$. Let $j:=\mathrm{argmin}_i u_i$ (ties broken arbitrarily). Assume that there is a $k\neq j$ with non-zero t_k *and* $u_k>u_j$. Define $\boldsymbol{t}'$ as above with $0<\varepsilon<\min\{t_k\,,t_k-t_j\}$. Concavity of f implies

$$F(\boldsymbol{u}')-F(\boldsymbol{u})=[f(u_j+\sigma\varepsilon)-f(u_j)]-[f(u_k)-f(u_k-\sigma\varepsilon)]>0,$$

which contradicts the maximality assumption of $\boldsymbol{t}$. Hence $t_i=0$ if $\underline{u_i}$ is not minimal ($\tilde{u}$). The previous paragraph constructed the unique solution $\boldsymbol{u}^{\overline{F}}$ satisfying this condition. Since this is the only local maximum it must be the unique global maximum (contrast this to the minimum case). ■

Theorem 2 (Exact extrema of expected entropy) *Let $\mathcal{H}(\boldsymbol{\pi})=-\sum_i\pi_i\log\pi_i$ be the entropy of $\boldsymbol{\pi}$ and the uncertainty of $\boldsymbol{\pi}$ be modeled by the Imprecise Dirichlet Model. The expected entropy $H(\boldsymbol{u}):=E_{\boldsymbol{t}}[\mathcal{H}]$ for given hyperparameter $\boldsymbol{t}$ and sample $\boldsymbol{n}$ is given by*

$$H(\boldsymbol{u})=\sum_i h(u_i) \quad \textit{with} \quad h(u)=u\cdot[\psi(n+s+1)-\psi((n+s)u+1)]=u\cdot\sum_{k=(n+s)u+1}^{n+s}k^{-1} \quad (5)$$

[6]Slope $\frac{f(u+\varepsilon)-f(u)}{\varepsilon}$ is a decreasing function in u for any $\varepsilon>0$, since f is concave.

where $\psi(x) = d\log\Gamma(x)/dx$ is the logarithmic derivative of the Gamma function and the last expression is valid for integral s and $(n+s)u$. The lower $\underline{H}$ and upper $\overline{H}$ expected entropies are assumed at $\boldsymbol{u}^{\underline{H}}$ and $\boldsymbol{u}^{\overline{H}}$ given in (3) and (4) (with F replaced by H, see also (1)).

A derivation of the exact expression (5) for the expected entropy can be found in [WW95, Hut01]. The only thing to be shown is that h is concave. This may be done by exploiting special properties of the digamma function ψ (see [AS74, Chp.6]). There are fast implementations of ψ and its derivatives and exact expressions for integer and half-integer arguments (see Appendix A for details).

Example 3 (Exact robust expected entropy) To see how the derived formulas can be used, let us compute the upper and lower expected entropy for for

$$d=2, \quad n_1=3, \quad n_2=6, \quad \text{i.e.} \quad n=9, \quad \text{and} \quad s=1, \quad \text{hence} \quad \sigma=\tfrac{1}{10}$$

The general correspondence (1) becomes

$$u_1=\tfrac{3+t_1}{10}, \quad u_2=\tfrac{6+t_2}{10}, \quad \text{hence} \quad \boldsymbol{t}^0=\boldsymbol{0} \quad \text{implies} \quad \boldsymbol{u}^0=\tbinom{0.3}{0.6}.$$

Using $n_1<n_2$, (3) implies

$$i^{\underline{H}}=2, \quad \boldsymbol{t}^{\underline{H}}=\tbinom{0}{1}, \quad \text{hence} \quad \boldsymbol{u}^{\underline{H}}=\tbinom{0.3}{0.7}.$$

From (4), using $i_1=1$ and $i_2=2$, we get

$$\tilde{u}=\min\left\{\tfrac{1+3}{9+1}, \tfrac{1+3+6}{2\cdot(9+1)}\right\}=\tfrac{4}{10}, \quad \text{hence} \quad \boldsymbol{u}^{\overline{H}}=\max\{\boldsymbol{u}^0,\tilde{u}\}=\tbinom{0.4}{0.6}.$$

This shows that the upper bound is assumed in a/the corner $\boldsymbol{t}^{\overline{H}}=\tbinom{1}{0}$. Inserting these u into (5), we get

$$h(\tfrac{3}{10})=\tfrac{2761}{8400}, \qquad h(\tfrac{4}{10})=\tfrac{2131}{6300}, \qquad h(\tfrac{6}{10})=\tfrac{1207}{4200}, \qquad h(\tfrac{7}{10})=\tfrac{847}{3600}.$$

Putting everything together we get the robust H estimate

$$\begin{aligned} \underline{\overline{H}} &= [H(\boldsymbol{u}^{\underline{H}}), H(\boldsymbol{u}^{\overline{H}})] = [h(\tfrac{3}{10})+h(\tfrac{7}{10})\,,\, h(\tfrac{4}{10})+h(\tfrac{6}{10})] \\ &= [\tfrac{7106}{12600}\,,\, \tfrac{7883}{12600}] \doteq [0.5639, 0.6256] \end{aligned}$$

The size of this interval is $\frac{37}{600}$, so $\overline{H}-\underline{H} \doteq 0.0616$ is of the order of σ. $\diamond$

In general, in order to apply Theorem 1, we need to be able to (a) somehow compute $F(\boldsymbol{u})$, e.g. compute the expectation $E_{\boldsymbol{t}}[\mathcal{F}]$, (b) verify whether $F(\boldsymbol{u})$ has the form $\sum_i f(u_i)$, which is often trivial, e.g. if $\mathcal{F}(\boldsymbol{\pi})=\sum_\rangle \phi(\pi_\rangle)$, and (c) prove concavity or convexity of F. In the following sections we derive conservative approximations for more general $F(\boldsymbol{u})$.

4 Approximate Robust Intervals

In this section we derive approximations for $\overline{\underline{F}}$ suitable for arbitrary, twice differentiable functions $F(\boldsymbol{u})$. The derived approximations for $\overline{\underline{F}}$ will be robust in the sense of covering set $\overline{\underline{F}}$ (for any n), and the approximations will be "good" if n is not too small. We do this by means of a finite Taylor series expansion in $\sigma:=\frac{s}{n+s}$ and by bounding the remainder.

In the following, we treat σ as a (small) expansion parameter. For $\boldsymbol{u},\boldsymbol{u}^*\in\Delta'$ we have

$$u_i-u_i^* \;=\; \sigma\cdot(t_i-t_i^*) \quad\text{and}\quad |u_i-u_i^*| \;=\; \sigma|t_i-t_i^*| \;\leq\; \sigma \quad\text{with}\quad \sigma:=\tfrac{s}{n+s}. \tag{6}$$

Hence we may Taylor-expand $F(\boldsymbol{u})$ around $\boldsymbol{u}^*$, which leads to a Taylor series in σ. This shows that F is approximately linear in $\boldsymbol{u}$ and hence in $\boldsymbol{t}$. A linear function on a simplex assumes its extreme values at the vertices of the simplex. This has already been encountered in Section 3. The consideration above is a simple explanation for this fact. This also shows that the robust interval $\overline{\underline{F}}$ is of size $\overline{F}-\underline{F}=O(\sigma)$.[7] Any approximation to $\overline{\underline{F}}$ should hence be at least $O(\sigma^2)$. The expansion of F to $O(\sigma)$ is

$$F(\boldsymbol{u}) \;=\; \overbrace{F(\boldsymbol{u}^*)}^{F_0=O(1)}+\overbrace{\sum_i[\partial_i F(\check{\boldsymbol{u}})](u_i-u_i^*)}^{F_R=O(\sigma)}, \tag{7}$$

where $\partial_i F(\check{\boldsymbol{u}})$ is the partial derivative $\partial F(\check{\boldsymbol{u}})/\partial\check{u}_i$ of $F(\check{\boldsymbol{u}})$ w.r.t. $\check{u}_i$. For suitable $\check{\boldsymbol{u}}=\check{\boldsymbol{u}}(\boldsymbol{u},\boldsymbol{u}^*)\in\Delta'$ this expansion is exact (F_R is the exact remainder). Natural points for expansion are $t_i^*=\frac{1}{d}$ in the center of Δ, or possibly also $t_i^*=\frac{n_i}{n}=u_i^*$. Here, we expand around the improper point $t_i^*:=t_i^0\equiv 0$, which is outside(!) Δ, since this makes expressions particularly simple.[8] Eq.(6) is still valid in this case, and F_R is exact for some $\check{\boldsymbol{u}}$ in

$$\Delta'_e:=\{\boldsymbol{u}\,:\,u_i\geq u_i^0\,\forall i,\;u_+\leq 1\},\quad\text{where}\quad u_i^0=\frac{n_i}{n+s}.$$

Note that we keep the exact condition $\boldsymbol{u}\in\Delta'$. F is usually already defined on Δ'_e or extends from Δ' to Δ'_e without effort in a natural way (analytical continuation). We introduce the notation

$$F\sqsubseteq G \qquad :\Leftrightarrow \qquad F\leq G \quad\text{and}\quad F=G+O(\sigma^2), \tag{8}$$

stating that G is a "good" upper bound on F. The following bounds hold for arbitrary differentiable functions. In order for the bounds to be "good," F has to be Lipschitz differentiable in the sense that there exists a constant c such that

$$|\partial_i F(\boldsymbol{u})|\leq c \quad\text{and}\quad |\partial_i F(\boldsymbol{u})-\partial_i F(\boldsymbol{u}')|\leq c|\boldsymbol{u}-\boldsymbol{u}'|$$

[7] $f(\boldsymbol{n},\boldsymbol{t},s)=O(\sigma^k)\;:\Leftrightarrow\;\exists c\,\forall\boldsymbol{n}\in I\!N_0^d,\boldsymbol{t}\in\Delta,s>0:\;|f(\boldsymbol{n},\boldsymbol{t},s)|\leq c\sigma^k$, where $\sigma=\frac{s}{n+s}$.

[8] The order of accuracy $O(\sigma^2)$ we will encounter is the same for all choices of $\boldsymbol{u}^*$. The concrete numerical errors differ of course. The choice $\boldsymbol{t}^*=\boldsymbol{0}$ can lead to $O(d)$ smaller F_R than the natural center point $\boldsymbol{t}^*=\frac{1}{d}$, but is more likely a factor $O(1)$ larger. The exact numerical values depend on the structure of F.

$$\forall\,\boldsymbol{u},\boldsymbol{u}'\in\Delta'_e \quad \text{and} \quad \forall\,i\in\{1,...,d\} \tag{9}$$

If F depends also on $\boldsymbol{n}$, e.g. via σ or $\boldsymbol{u}^0$, then c shall be independent of them.

The Lipschitz condition is satisfied, for instance, if the curvature $\partial^2 F$ is uniformly bounded. This is satisfied for the expected entropy H (see (5)), but violated for the approximation $E_{\boldsymbol{t}}[\mathcal{H}]\approx\mathcal{H}(\boldsymbol{u})$ if $n_i=0$ for some i.

Theorem 4 (Approximate robust intervals) *Assume $F:\Delta'_e\to I\!R$ is a Lipschitz differentiable function (9). Let $[\underline{F},\overline{F}]$ be the global [minimum,maximum] of F restricted to Δ'. Then*

$$F(\boldsymbol{u}^1)\ \sqsubseteq\ \overline{F}\ \sqsubseteq\ F_0+F_R^{ub}\ \ \textit{where}\ \ F_R^{ub}:=\max_i F_{iR}^{ub}\ \textit{and}\ \ F_{iR}^{ub}\ :=\ \sigma\max_{\boldsymbol{u}\in\Delta'_e}[\partial_i F(\boldsymbol{u})],$$

$$F_0+F_R^{lb}\ \sqsubseteq\ \underline{F}\ \sqsubseteq\ F(\boldsymbol{u}^2)\ \ \textit{where}\ \ F_R^{lb}:=\min_i F_{iR}^{lb}\ \ \textit{and}\ \ F_{iR}^{lb}\ :=\ \sigma\min_{\boldsymbol{u}\in\Delta'_e}[\partial_i F(\boldsymbol{u})],$$

$F_0:=F(\boldsymbol{u}^0)$, and $t_i^1:=\delta_{ii^1}$ with $i^1:=\text{argmax}_i F_{iR}^{ub}$, and $t_i^2:=\delta_{ii^2}$ with $i^2:=\text{argmin}_i F_{iR}^{lb}$, and $\sqsubseteq$ defined in (8) means $\leq$ and $=$ $+O(\sigma^2)$, where $\sigma=1-u_+^0$.

For conservative estimates, the lower bound on $\underline{F}$ and the upper bound on $\overline{F}$ are the interesting ones. Together with the "inner" bounds $F(\boldsymbol{u}^1)$ and $F(\boldsymbol{u}^2)$, they also yield interesting information about the accuracy of the approximations: $F_0+F_R^{ub}-F(\boldsymbol{u}^1)$ is an upper bound on the (unknown) approximation error $F_0+F_R^{ub}-\overline{F}$, and similarly for $\underline{F}$.

Proof. We start by giving an $O(\sigma^2)$ bound on $\overline{F}_R=\max_{\boldsymbol{u}\in\Delta'}F_R(\boldsymbol{u})$. We first insert (6) with $\boldsymbol{t}^*=\boldsymbol{t}^0\equiv\boldsymbol{0}$ into (7) and treat $\check{\boldsymbol{u}}$ and $\boldsymbol{t}$ as separate variables:

$$F_R(\check{\boldsymbol{u}},\boldsymbol{t})\ =\ \sigma\sum_i[\partial_i F(\check{\boldsymbol{u}})]\cdot t_i\ \sqsubseteq\ \max_{\check{\boldsymbol{u}}\in\Delta'_e}\left\{\sigma\sum_i[\partial_i F(\check{\boldsymbol{u}})]\cdot t_i\right\}\ \sqsubseteq\ \sum_i F_{iR}^{ub}\cdot t_i$$

$$\text{with}\quad F_{iR}^{ub}\ :=\ \sigma\max_{\check{\boldsymbol{u}}\in\Delta'_e}[\partial_i F(\check{\boldsymbol{u}})] \tag{10}$$

The first inequality is obvious, the second follows from the convexity of max. From assumption (9) we get $\partial_i F(\boldsymbol{u})-\partial_i F(\boldsymbol{u}')=O(\sigma)$ for all $\boldsymbol{u},\boldsymbol{u}'\in\Delta'_e$, since Δ'_e has diameter $O(\sigma)$. Due to one additional σ in (10) the expressions in (10) change only by $O(\sigma^2)$ when introducing or dropping $\max_{\check{\boldsymbol{u}}}$ anywhere. This shows that the inequalities are tight within $O(\sigma^2)$ and justifies $\sqsubseteq$. We now upper bound $F_R(\boldsymbol{u})$:

$$\overline{F}_R=\max_{\boldsymbol{u}\in\Delta'}F_R(\boldsymbol{u})\sqsubseteq\max_{\boldsymbol{t}\in\Delta}\max_{\check{\boldsymbol{u}}\in\Delta'_e}F_R(\check{\boldsymbol{u}},\boldsymbol{t})\sqsubseteq\max_{\boldsymbol{t}\in\Delta}\sum_i F_{iR}^{ub}\cdot t_i=\max_i F_{iR}^{ub}=:F_R^{ub} \tag{11}$$

A linear function on Δ is maximized by setting the t_i component with largest coefficient to 1. This shows the last equality. The maximization over $\check{\boldsymbol{u}}$ in (10) can often be performed analytically, leaving an easy $O(d)$ time task for maximizing over i.

We have derived an upper bound F_R^{ub} on $\overline{F}_R$. Let us define the corner $t_i=\delta_{ii^1}$ of Δ with $i^1:=\text{argmax}_i F_{iR}^{ub}$. Since $\overline{F}_R\geq F_R(\boldsymbol{u})$ for all $\boldsymbol{u}$, $F_R(\boldsymbol{u}^1)$ in particular is a lower bound on $\overline{F}_R$. A similar line of reasoning as above shows that that $F_R(\boldsymbol{u}^1) = \overline{F}_R + O(\sigma^2)$. Using $\overline{F+const.} = \overline{F}+const.$ we get $O(\sigma^2)$ lower and upper bounds on $\overline{F}$, i.e. $F(\boldsymbol{u}^1) \sqsubseteq \overline{F} \sqsubseteq F_0 + F_R^{ub}$. $\underline{F}$ is bound similarly with all max's replaced by min's and inequalities reversed. Together this proves the Theorem 4. ■

In the following sections we assume the definitions/notation of Theorem 4 for F and analogous ones for all other occurring estimators (G,H,I,...).

5 Error Propagation

We now show how bounds of elementary functions obtained by Theorem 4 can be used to get bounds for more complex composite functions, especially for sums and products of functions. The results are used in Section 6 for deriving robust intervals for the mutual information for which exact solutions are not known.

Approximation of $\overline{\underline{F}}$ (special cases). For the special case $F(\boldsymbol{u}) = \sum_i f(u_i)$ we have $\partial_i F(\boldsymbol{u}) = f'(u_i)$. For concave f like in case of the entropy we get particularly simple bounds

$$\begin{aligned} F_{iR}^{ub} &= \sigma \max\nolimits_{\boldsymbol{u}\in\Delta'_e} f'(u_i) = \sigma f'(u_i^0), & F_R^{ub} &= \sigma \max_i f'(u_i^0) = \sigma f'(\tfrac{\min_i n_i}{n+s}), \\ F_{iR}^{lb} &= \sigma \min\nolimits_{\boldsymbol{u}\in\Delta'_e} f'(u_i) = \sigma f'(u_i^0+\sigma), & F_R^{lb} &= \sigma \min_i f'(u_i^0+\sigma) = \sigma f'(\tfrac{\max_i n_i+s}{n+s}), \end{aligned} \tag{12}$$

where we have used $\max_{\boldsymbol{u}\in\Delta'_e} f'(u_i) = \max_{u_i\in[u_i^0,u_i^0+\sigma]} f'(u_i) = f'(u_i^0)$, and similarly for min. Analogous results hold for convex functions. In case the maximum cannot be found exactly one is allowed to further increase Δ'_e as long as its diameter remains $O(\sigma)$. Often an increase to $\Box' := \{\boldsymbol{u} : u_i^0 \leq u_i \leq u_i^0+\sigma\} \supset \Delta'_e \supset \Delta'$ makes the problem easy. Note that if we were to perform these kind of crude enlargements on $\max_{\boldsymbol{u}} F(\boldsymbol{u})$ directly we would loose the bounds by $O(\sigma)$.

Example 5 (Approximate robust expected entropy) Let us compare the exact robust estimate of the expected entropy for $n_1=3$, $n_2=6$, $s=1$ (hence $n=9$, and $\sigma=\frac{1}{10}$) computed in Example 3 with this approximation: Using the expressions for h' from Appendix A, we get

$$h'(\tfrac{3}{10}) = \tfrac{13051}{2520} - \tfrac{1}{2}\Pi^2 \quad \text{and} \quad h'(\tfrac{7}{10}) = \tfrac{91717}{8400} - \tfrac{7}{6}\Pi^2,$$

where $\Pi \doteq 3.1415$. From (2) and (12) we get

$$H_0 = H(\boldsymbol{u}^0) = h(\tfrac{3}{10}) + h(\tfrac{6}{10}) = \tfrac{69}{112}, \qquad H_R^{ub} = \tfrac{1}{10}h'(\tfrac{3}{10}), \qquad H_R^{lb} = \tfrac{1}{10}h'(\tfrac{7}{10}).$$

Together with the expressions from Example 3 we get the conservative estimate

$$[H_0 + H_R^{lb}\,,\, H_0 + H_R^{ub}] \doteq [0.5564, 0.6404].$$

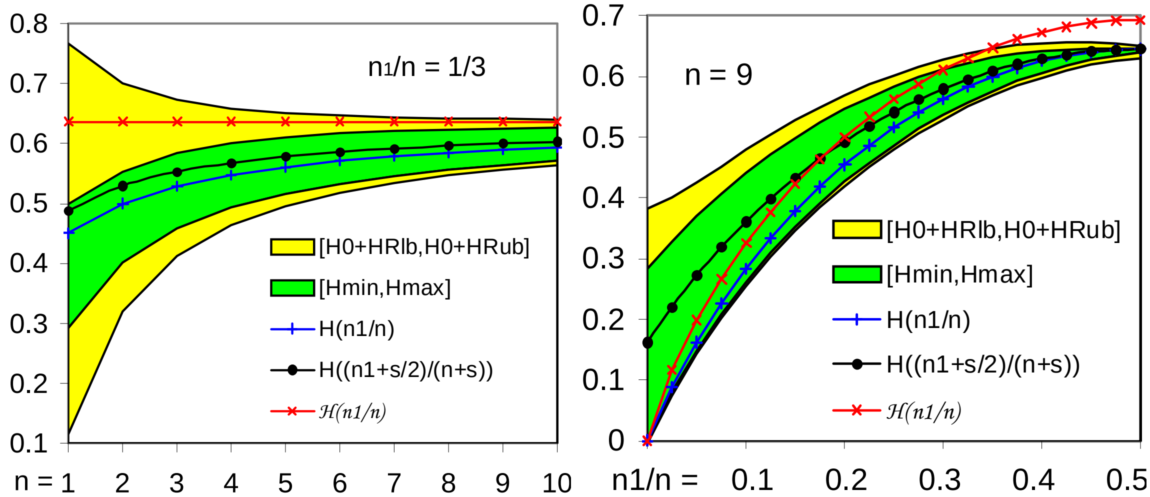

Figure 1: [Expected Entropy] The figures display the various (expected) entropy estimates for $s=1$: The left figure for $n_1/n=1/3$ and $n=1...10$. The right figure for $n=9$ and $n_1/n=0...0.5$. The "intersection" $n_1=3$ and $n_2=6$ is treated analytically in Examples 3 and 5. The green (dark gray) area is the exact robust interval $[\underline{H}\,,\overline{H}]$ from Theorem 2. The yellow+green (gray) area is the conservative estimate $[H_0+H_R^{lb}\,,H_0+H_R^{ub}]$ from Theorem 4. The area $[H(\boldsymbol{u}^2)\,,H(\boldsymbol{u}^1)]$ is not shown, since (here) it essentially coincides with $\overline{\underline{H}}$). Some point estimates $H(\frac{\boldsymbol{n}}{n})$, $H(\frac{\boldsymbol{n}+\boldsymbol{1}/2}{n+1})$, and $\mathcal{H}(\frac{\boldsymbol{n}}{n})$ are also shown.

The approximation accuracy

$$H_0+H_R^{ub}-\overline{H} \doteq 0.0148 \quad \text{and} \quad \underline{H}-H_0-H_R^{lb} \doteq 0.0074$$

is consistent with our $O(\sigma^2)$ estimation. If exact expressions are not available we can upper bound the widening by

$$H_0+H_R^{ub}-H(\boldsymbol{u}^1) \doteq 0.0148 \quad \text{and} \quad H(\boldsymbol{u}^2)-H_0-H_R^{lb} \doteq 0.0074$$

Since generally $\boldsymbol{u}^2=\boldsymbol{u}^{\underline{H}}$ and in our example also $\boldsymbol{u}^1=\boldsymbol{u}^{\overline{H}}$, the numbers coincide. $\diamondsuit$

Example 6 (Entropy: dependency on $\boldsymbol{n}$) Figure 1 (left) shows how the size of the (conservative) robust interval of the expected entropy H varies with the sample size n. We considered $s=1$ and $d=2$ and kept $n_1/n={}^1\!/_3$ and $n_2/n={}^2\!/_3$ fix (allowing for fractional $\boldsymbol{n}$). We clearly see that the yellow (light gray) region diminishes quickly compared to the green (dark gray) region with increasing n, i.e. the approximation accuracy gets better for larger n. Some point estimates $H(\frac{\boldsymbol{n}}{n})$, $H(\frac{\boldsymbol{n}+\boldsymbol{1}/2}{n+1})$, and $\mathcal{H}(\frac{\boldsymbol{n}}{n})$ are also shown. Figure 1 (right) shows the intervals for fixed $n=9$, while varying $n_1/n=0...0.5$ ($n_1/n=0.5...1$ is symmetric). The interval $\overline{\underline{H}}$ is shorter for more uniform $\boldsymbol{u}$, since H (like $\mathcal{H}$) varies more closer to the boundary of

Δ. The $[H(\boldsymbol{u}^2),H(\boldsymbol{u}^1)]$ region is not shown since it is identical to $\underline{\overline{H}}$ (also in the left graph except for $n=1$). For $n=9$ and $n_1/n=1/3$ we recover the results of Examples 3 and 5 (left and right figure). $\diamondsuit$

Error propagation. Assume we found bounds for estimators $G(\boldsymbol{u})$ and $H(\boldsymbol{u})$ and we want now to bound the sum $F(\boldsymbol{u}):=G(\boldsymbol{u})+H(\boldsymbol{u})$. In the direct approach $\overline{F}\leq\overline{G}+\overline{H}$ we may lose $O(\sigma)$. A simple example is $G(\boldsymbol{u})=u_i$ and $H(\boldsymbol{u})=-u_i$ for which $F(\boldsymbol{u})=0$, hence $0=\overline{F}\leq\overline{G}+\overline{H}=u_i^0+\sigma-u_i^0=\sigma$, i.e. $\overline{F}\not\sqsupseteq\overline{G}+\overline{H}$. We can exploit the techniques of the previous section to obtain $O(\sigma^2)$ approximations.

$$F_{iR}^{ub} \;=\; \sigma\max_{\boldsymbol{u}\in\Delta'_e}\partial_i F(\boldsymbol{u}) \;\sqsubseteq\; \sigma\max_{\boldsymbol{u}\in\Delta'_e}\partial_i G(\boldsymbol{u})+\sigma\max_{\boldsymbol{u}\in\Delta'_e}\partial_i H(\boldsymbol{u}) \;=\; G_{iR}^{ub}+H_{iR}^{ub}$$

Theorem 7 (Error propagation: Sum) *Let $G(\boldsymbol{u})$ and $H(\boldsymbol{u})$ be Lipschitz differentiable and $F(\boldsymbol{u})=\alpha G(\boldsymbol{u})+\beta H(\boldsymbol{u})$, $\alpha,\beta\geq 0$, then $\overline{F}\sqsubseteq F_0+F_R^{ub}$ and $\underline{F}\sqsupseteq F_0+F_R^{lb}$, where $F_0=\alpha G_0+\beta H_0$, and $F_{iR}^{ub}\sqsubseteq\alpha G_{iR}^{ub}+\beta H_{iR}^{ub}$, and $F_{iR}^{lb}\sqsupseteq\alpha G_{iR}^{lb}+\beta H_{iR}^{lb}$.*

It is important to notice that $F_R^{ub}\not\sqsupseteq G_R^{ub}+H_R^{ub}$ (use previous example), i.e. $\max_i[G_{iR}^{ub}+H_{iR}^{ub}]\not\sqsupseteq\max_i G_{iR}^{ub}+\max_i H_{iR}^{ub}$. $\max_i$ can not be pulled in and it is important to propagate F_{iR}^{ub}, rather than F_R^{ub}.

Every function F with bounded curvature can be written as a sum of a concave function G and a convex function H. For convex and concave functions, determining bounds is particularly easy, as we have seen. Often F decomposes naturally into convex and concave parts as is the case for the mutual information, addressed later. Bounds can also be derived for products.

Theorem 8 (Error propagation: Product) *Let $G,H:\Delta'_e\to[0,\infty)$ be non-negative Lipschitz differentiable functions (9) with non-negative derivatives $\partial_i G,\partial_i H\geq 0$ $\forall i$ and $F(\boldsymbol{u})=G(\boldsymbol{u})\cdot H(\boldsymbol{u})$, then $\overline{F}\sqsubseteq F_0+F_R^{ub}$, where $F_0=G_0\cdot H_0$, and $F_{iR}^{ub}\sqsubseteq G_{iR}^{ub}(H_0+H_R^{ub})+(G_0+G_R^{ub})H_{iR}^{ub}$, and similarly for $\underline{F}$.*

Proof. We have

$$\begin{aligned}F_{iR}^{ub} \;&=\; \sigma\max\partial_i F \;=\; \sigma\max\partial_i(G\cdot H) \;=\; \sigma\max[(\partial_i G)H+G(\partial_i H)] \;\sqsubseteq \\ \sigma(\max\partial_i G)&(\max H)+\sigma(\max G)(\max\partial_i H)\sqsubseteq G_{iR}^{ub}(H_0+H_R^{ub})+(G_0+G_R^{ub})H_{iR}^{ub}\end{aligned}$$

where all functions depend on $\boldsymbol{u}$ and all max are over $\boldsymbol{u}\in\Delta'_e$. There is one subtlety in the last inequality: $\max G\neq\overline{G}\sqsubseteq G_0+G_R^{ub}$. The reason for the $\neq$ being that the maximization is taken over Δ'_e, not over Δ' as in the definition of $\overline{G}$. The correct line of reasoning is as follows:

$$\max_{\boldsymbol{u}\in\Delta'_e}G_R(\boldsymbol{u})\sqsubseteq\max_{\boldsymbol{t}\in\Delta_e}\sum_i G_{iR}^{ub}\cdot t_i=\max\{0,\max_i G_{iR}^{ub}\}=G_R^{ub}\;\Rightarrow\;\max G\sqsubseteq G_0+G_R^{ub}$$

The first inequality can be proven in the same way as (11). In the first equality we set the $t_i=1$ with maximal G_{iR}^{ub} *if* it is positive. If all G_{iR}^{ub} are negative we set $\boldsymbol{t}\equiv\boldsymbol{0}$.

We assumed $G\geq 0$ and $\partial_i G\geq 0$, which implies $G_R\geq 0$. So, since $G_R\geq 0$ anyway, this subtlety is ineffective. Similarly for $\max H_R$. ■

It is possible to remove the rather strong non-negativity assumptions. Propagation of errors for other combinations like ratios $F=G/H$ may also be obtained.

6 Robust Intervals for Expected Mutual Information

We illustrate the application of the previous results on the Mutual Information between two random variables $\imath\in\{1,...,d_1\}$ and $\jmath\in\{1,...,d_2\}$.

Mutual Information. Consider an i.i.d. random process with outcome $(i,j)\in\{1,...,d_1\}\times\{1,...,d_2\}$ having joint probability π_{ij}, where $\boldsymbol{\pi}\in\Delta:=\{\boldsymbol{x}\in I\!R^{d_1\times d_2}:x_{ij}\geq 0\,\forall ij,\, x_{++}=1\}$. An important measure of the stochastic dependence of $\imath$ and $\jmath$ is the mutual information

$$
\begin{aligned}
\mathcal{I}(\boldsymbol{\pi}) &= \sum_{i=1}^{d_1}\sum_{j=1}^{d_2}\pi_{ij}\log\frac{\pi_{ij}}{\pi_{i+}\pi_{+j}} &&= \sum_{ij}\pi_{ij}\log\pi_{ij}-\sum_i\pi_{i+}\log\pi_{i+}-\sum_j\pi_{+j}\log\pi_{+j} \\
& &&= \mathcal{H}(\boldsymbol{\pi}_{\imath+})+\mathcal{H}(\boldsymbol{\pi}_{+\jmath})-\mathcal{H}(\boldsymbol{\pi}_{\imath\jmath}),
\end{aligned} \tag{13}
$$

where $\pi_{i+}=\sum_j\pi_{ij}$ and $\pi_{+j}=\sum_i\pi_{ij}$ are row and column marginal chances. Again, we assume a Dirichlet prior over $\boldsymbol{\pi}_{\imath\jmath}$, which leads to a Dirichlet posterior $p(\boldsymbol{\pi}_{\imath\jmath}|\boldsymbol{n})\propto\prod_{ij}\pi_{ij}^{n_{ij}+st_{ij}-1}$ with $\boldsymbol{t}\in\Delta$. The expected value of π_{ij} is

$$E_{\boldsymbol{t}}[\pi_{ij}] = \frac{n_{ij}+st_{ij}}{n+s} =: u_{ij}$$

The marginals $\boldsymbol{\pi}_{i+}$ and $\boldsymbol{\pi}_{+j}$ are also Dirichlet with expectation u_{i+} and u_{+j}. The expected mutual information $I(\boldsymbol{u}):=E_{\boldsymbol{t}}[\mathcal{I}]$ can, hence, be expressed in terms of the expectations of three entropies $H(\boldsymbol{u}):=E_{\boldsymbol{t}}[\mathcal{H}]$ (see (5))

$$
\begin{aligned}
I(\boldsymbol{u}) &= H(\boldsymbol{u}_{\imath+})+H(\boldsymbol{u}_{+\jmath})-H(\boldsymbol{u}_{\imath\jmath}) \;=\; H_{row}+H_{col}-H_{joint} \\
&= \sum_i h(u_{i+})+\sum_j h(u_{+j})-\sum_{ij}h(u_{ij}),
\end{aligned}
$$

where here and in the following we index quantities with *joint*, *row*, and *col* to denote to which distribution the quantity refers.

Crude bounds for $\boldsymbol{I(u)}$. Estimates for the robust IDM interval $[\min_{\boldsymbol{t}\in\Delta}E_{\boldsymbol{t}}[\mathcal{I}]\,,\max_{\boldsymbol{t}\in\Delta}E_{\boldsymbol{t}}[\mathcal{I}]]$ can be obtained by [minimizing,maximizing] $I(\boldsymbol{u})$. A crude upper bound can be obtained as

$$
\begin{aligned}
\overline{I} &:= \max_{\boldsymbol{t}\in\Delta}I(\boldsymbol{u}) \;=\; \max[H_{row}+H_{col}-H_{joint}] \;\leq \\
&\quad \max H_{row}+\max H_{col}-\min H_{joint} \;=\; \overline{H}_{row}+\overline{H}_{col}-\underline{H}_{joint},
\end{aligned}
$$

where exact solutions to $\overline{H}_{row}$, $\overline{H}_{col}$ and $\underline{H}_{joint}$ are available from Section 3. Similarly $\underline{I}\geq\underline{H}_{row}+\underline{H}_{col}-\overline{H}_{joint}$. The problem with these bounds is that, although good in some cases, they can become arbitrarily crude. The following $O(\sigma^2)$ bound can be derived by exploiting the error sum propagation Theorem 7.

Theorem 9 (Bound on lower and upper expected Mutual Information) *The following bounds on the expected mutual information $I(\boldsymbol{u})=E_{\boldsymbol{t}}[\mathcal{I}]$ are valid:*

$$
\begin{aligned}
&I(\boldsymbol{u}^1)\sqsubseteq\overline{I}\sqsubseteq I_0+I_R^{ub} \quad \textit{and} \quad I_0+I_R^{lb}\sqsubseteq\underline{I}\sqsubseteq I(\boldsymbol{u}^2), \quad \textit{where}\\
&I_0=I(\boldsymbol{u}^0)=H_{0row}+H_{0col}-H_{0joint}=\textstyle\sum_i h(u_{i+}^0)+\sum_j h(u_{+j}^0)-\sum_{ij}h(u_{ij}^0),\\
&I_{ijR}^{ub}\sqsubseteq H_{iRrow}^{ub}+H_{jRcol}^{ub}-H_{ijRjoint}^{lb}=h'(u_{i+}^0)+h'(u_{+j}^0)-h'(u_{ij}^0+\sigma),\\
&I_{ijR}^{lb}\sqsupseteq H_{iRrow}^{lb}+H_{jRcol}^{lb}-H_{ijRjoint}^{ub}=h'(u_{i+}^0+\sigma)+h'(u_{+j}^0+\sigma)-h'(u_{ij}^0),
\end{aligned}
$$

with h defined in (5), and $t_{ij}^0=0$, and $t_{ij}^1=\delta_{(ij)(ij)^1}$ with $(ij)^1=\mathrm{argmax}_{ij}I_{ijR}^{ub}$, and $t_{ij}^2=\delta_{(ij)(ij)^2}$ with $(ij)^2=\mathrm{argmin}_{ij}I_{ijR}^{lb}$, and $I_R^{ub}=\max_{ij}I_{ijR}^{ub}$, and $I_R^{lb}=\max_{ij}I_{ijR}^{lb}$.

7 The IDM for Product Spaces

In the last section we considered the "full" IDM on the product of two random variables. The structure of the problem suggests considering a smaller "product" of IDMs as described below, which can lead to better estimates.

Product spaces $\Omega=\Omega_1\times...\times\Omega_m$ with $\Omega_k=\{1,...d_k\}$ occur frequently in practical problems, e.g. in the mutual information ($m=2$), in robust trees ($m=3$), or in Bayesian nets in general (m large). Without loss of generality we only discuss the $m=2$ case in the following. Ignoring the underlying structure in Ω, a Dirichlet prior in case of unknown chances π_{ij} and an IDM as used in Section 6 with

$$\boldsymbol{t}\in\Delta:=\{\boldsymbol{t}\in I\!R^{d_1\times d_2}\equiv I\!R^{d_1}\otimes I\!R^{d_2}\;:\;t_{ij}\geq 0\,\forall ij,\;t_{++}=1\} \tag{14}$$

seems natural.

On the other hand, if we take into account the structure of Ω and go back to the original motivation of the IDM, this choice is far less obvious. Recall that one of the major motivations of the IDM was its representation invariance in the sense that inferences are not affected when grouping or splitting events in Ω. For unstructured spaces like Ω_k this is a reasonable principle. For illustration, let us consider objects of various *shape* and *color*, i.e. $\Omega=\Omega_1\times\Omega_2$, $\Omega_1=\{ball,pen,die,...\}$, $\Omega_2=\{yellow,red,green,...\}$ in generalization to Walley's bag of marbles example. Assume we want to detect a potential dependency between *shape* and *color* by means of their mutual information I. If we have no prior idea on the possible kind of colors, a model which is independent of the choice of Ω_2 is welcome. Grouping red and green, for instance, corresponds to grouping (x_{i1}, x_{i2}, x_{i3}, x_{i4},...) to (x_{i1}, $x_{i2}+x_{i3}$, x_{i4},...) *for all shapes i*, where $\boldsymbol{x}\in\{\boldsymbol{n},\boldsymbol{\pi},\boldsymbol{t},\boldsymbol{u}\}$. Similarly for the different shapes, for instance we could group all round or all angular objects. The "smallest IDM" which respects this invariance is the one which considers all

$$\boldsymbol{t} \in \mathbb{\Delta} := \Delta_{d_1} \otimes \Delta_{d_2} \; \subsetneq \; \Delta. \tag{15}$$

The tensor or outer product $\otimes$ is defined as $(\boldsymbol{v}\otimes\boldsymbol{w})_{ij} := v_i w_j$ and $V \otimes W := \{\boldsymbol{v}\otimes\boldsymbol{w} : \boldsymbol{v}\in V, \boldsymbol{w}\in W\}$. It is a bilinear (not linear!) mapping. This smaller product IDM $\mathbb{\Delta}$ is invariant under arbitrary grouping of columns and rows of the chance matrix $(\boldsymbol{\pi}_{ij})_{1\leq i\leq d_1, 1\leq j\leq d_2}$. In contrast to the larger full IDM Δ it is not invariant under arbitrary grouping of matrix cells, but there is anyway little motivation for the necessity of such a general invariance. General non-column/row cross groupings would destroy the product structure of Ω and with that the mere concepts of shape and color, and their correlation. For $m>2$ as in Bayes-nets cross groupings look even less natural. Whether the $\mathbb{\Delta}$ or the larger simplex Δ is the more appropriate IDM depends on whether one regards the structure $\Omega_1\times\Omega_2$ of Ω as a natural prior knowledge or as an arbitrary a posteriori choice. The smaller IDM has the potential advantage of leading to more precise predictions (smaller robust sets).

Let us consider an estimator $F:\Delta\to I\!\!R$ and its restriction $F_\otimes:\mathbb{\Delta}\to I\!\!R$. Robust intervals $[\underline{F},\overline{F}]$ for Δ are generally wider than robust intervals $[\underline{F}_\otimes,\overline{F}_\otimes]$ for $\mathbb{\Delta}$. Fortunately not much. Although $\mathbb{\Delta}$ is a *lower-dimensional* subspace of Δ, it contains all vertices of Δ. This is possible since $\mathbb{\Delta}$ is a *nonlinear* subspace. The set of "vertices" in both cases is $\{\boldsymbol{t} : t_{ij} = \delta_{ii_0}\delta_{jj_0},\, i_0\in\Omega_1,\, j_0\in\Omega_2\}$. Hence, *if* the robust interval boundaries $\overline{F}$ are assumed in the vertices of Δ *then* the interval for the $\mathbb{\Delta}$ IDM model is the same ($\overline{F} = \overline{F}_\otimes$). Since the condition is "approximately" true, the conclusion is "approximately" true. More precisely:

Theorem 10 (IDM bounds for product spaces) *The $O(\sigma^2)$ bounds of Theorem 4 on the robust interval $\overline{F}$ in the full IDM Δ (14), remain valid for $\overline{F}_\otimes$ in the product IDM $\mathbb{\Delta}$ (15).*

Proof.

$$F(\boldsymbol{u}^1) \leq \overline{F}_\otimes \leq \overline{F} \leq F_0 + F_R^{ub} = F(\boldsymbol{u}^1) + O(\sigma^2),$$

where $\overline{F}_\otimes := \max_{\boldsymbol{t}\in\mathbb{\Delta}} F(\boldsymbol{u})$ and $\boldsymbol{u}^1$ was the "F_R maximizing" vertex as defined in Theorem 9 ($F(\boldsymbol{u}^1) \sqsubseteq \overline{F}$). The first inequality follows from the fact that all Δ vertices also belong to $\mathbb{\Delta}$, i.e. $\boldsymbol{t}^1\in\mathbb{\Delta}$. The second inequality follows from $\mathbb{\Delta}\subset\Delta$. The remaining (in)equalities follow from Theorem 4. This shows that $|\overline{F}_\otimes - \overline{F}| = O(\sigma^2)$, hence $F_0 + F_R^{ub}$ is also an $O(\sigma^2)$ upper bound to $\overline{F}_\otimes$. This implies that to the approximation accuracy we can achieve, the choice between Δ and $\mathbb{\Delta}$ is irrelevant. ∎

8 Robust Credible Intervals

So far we have considered robust intervals of *expected* values $F=E_{\boldsymbol{t}}[\mathcal{F}]$. We now briefly consider the problem of how to combine Bayesian credible intervals for $\mathcal{F}$ with robust intervals of the IDM.

Bayesian credible sets/intervals. For a probability density $p:I\!R^d\to[0,1]$, an α-credible region is a measurable set A for which $p(A):=\int p(x)1\!\!1_A(x)d^dx\geq\alpha$, where $1\!\!1_A(x)=1$ if $x\in A$ and 0 otherwise, i.e. $x\in A$ with probability at least α. For given α, there are many choices for A. Often one is interested in "small" sets, where the size of A may be measured by its volume $\text{Vol}(A):=\int 1\!\!1_A(x)d^dx$. Let us define a/the smallest α-credible set

$$A^{min}:=\mathop{\arg\min}_{A:p(A)\geq\alpha}\text{Vol}(A)$$

with ties broken arbitrarily. For unimodal p, A^{min} can be chosen as a connected set. For $d=1$ this means that $A^{min}=[a,b]$ with $\int_a^b p(x)dx=\alpha$ is a minimal length highest density α-credible interval. If, additionally p is symmetric around $E[x]$, then $A^{min}=[E[x]-c,E[x]+c]$ is also symmetric around $E[x]$.

Robust credible sets. If we have a set of probability distributions $\{p_t(x),\,t\in T\}$, we can choose for each t an α-credible set A_t with $p_t(A_t)\geq\alpha$, a minimal one being $A_t^{min}:=\text{argmin}_{A:p_t(A)\geq\alpha}\text{Vol}(A)$. A robust α-credible set is a set A which contains x with p_t-probability at least α for *all* t. A minimal size robust α-credible set is

$$A^{min}:=\mathop{\arg\min}_{A=\bigcup_t A_t:p_t(A_t)\geq\alpha}\text{Vol}(A). \tag{16}$$

It is not easy to deal with this expression, since A^{min} is *not* a function of $\{A_t^{min}:t\in T\}$, and especially does not coincide with $\bigcup_t A_t^{min}$ as one might expect.

Robust credible intervals. This can most easily be seen for univariate symmetric unimodal distributions, where t is a translation, e.g. $p_t(x)=\text{Normal}(E_t[x]=t,\sigma=1)$ with 95% credible intervals $A_t^{min}=[t-2,t+2]$. For, e.g. $T=[-1,1]$ we get $\bigcup_t A_t^{min}=[-3,3]$. The credible intervals *move* with t. One can get a smaller union if we take the intervals $A_t^{sym}=[-c_t,c_t]$ symmetric around 0. Since A_t^{sym} is a non-central interval w.r.t. p_t for $t\neq 0$, we have $c_t>2$, i.e. A_t^{sym} is larger than A_t^{min}, but one can show that the increase of c_t is smaller than the shift of A_t^{min} by t, hence we save something in the union. The optimal choice is neither A_t^{sym} nor A_t^{min}, but something in-between.

To illustrate this point numerically consider triangular distributions instead of Gaussians:

$$p_t(x):=\max\{0\,,\,1-|x-t|\},\qquad t\in T:=[-\gamma,\gamma],\qquad \gamma>0,$$

$$\Rightarrow\; p_t([a,b])=\left|b^*(1-\tfrac{1}{2}|b^*|)-a^*(1-\tfrac{1}{2}|a^*|)\right| \text{ with } \begin{array}{l} a^*=\min\{\max\{a,0\},1\}-t, \\ b^*=\min\{\max\{b,0\},1\}-t. \end{array}$$

One can derive the following expressions for the α-credible intervals, valid for (the interesting case of) $\alpha\geq\frac{1}{2}$.

$$A_t^{min} = [t-1+\sqrt{1-\alpha}\,,\, t+1-\sqrt{1-\alpha}],$$
$$\bigcup_{t\in T} A_t^{min} = [-\gamma-1+\sqrt{1-\alpha}\,,\, \gamma+1-\sqrt{1-\alpha}].$$
$$A^{min} = \left\{ \begin{array}{cl} \Big[-1+\sqrt{1-\alpha-\gamma^2}\,,\, 1-\sqrt{1-\alpha-\gamma^2}\Big] & \text{for} \quad \gamma^2 \leq \frac{1}{2}(1-\alpha), \\ \Big[-\gamma-1+\sqrt{2(1-\alpha)}\,,\, \gamma+1-\sqrt{2(1-\alpha)}\Big] & \text{for} \quad \gamma^2 \geq \frac{1}{2}(1-\alpha). \end{array} \right.$$

It is easy to see that $A^{min}\subset\bigcup_t A_t^{min}$ and that A^{min} is a proper subinterval of $\bigcup_t A_t^{min}$ of shorter length for every $\gamma>0$ and $\frac{1}{2}\leq\alpha<1$.

An interesting open question is under which general conditions we can expect $A^{min}\subseteq\bigcup_t A_t^{min}$. In any case, $\bigcup_t A_t$ can be used as a conservative estimate for a robust credible set, since $p_t(\bigcup_{t'} A_{t'})\geq p_t(A_t)\geq\alpha$ for all t.

A special (but important) case which falls outside the above framework are one-sided credible intervals, where only A_t of the form $[a,\infty)$ are considered. In this case $A^{min}=\bigcup_t A_t^{min}$, i.e. $A^{min}=[a_{min},\infty)$ with $a_{min}=\max\{a: p_t([a,\infty])\geq\alpha\forall t\}$.

Approximations. For complex distributions like for the mutual information we have to approximate (16) somehow. We use the following notation for shortest α-credible *intervals* w.r.t. a univariate distribution $p_t(x)$:

$$\underset{\sim}{\widetilde{x}_t} \;\equiv\; [\underset{\sim}{x_t}, \widetilde{x}_t] \;\equiv\; [E_t[x]-\Delta\underset{\sim}{x_t}\,,\, E_t[x]+\Delta\widetilde{x}_t] \;:=\; \mathop{\arg\min}_{[a,b]:p_t([a,b])\geq\alpha}(b-a),$$

where $\Delta\widetilde{x}_t := \widetilde{x}_t - E_t[x]$ ($\Delta\underset{\sim}{x_t} := E_t[x]-\underset{\sim}{x_t}$) is the distance from the right boundary $\widetilde{x}_t$ (left boundary $\underset{\sim}{x_t}$) of the shortest α-credible interval $\underset{\sim}{\widetilde{x}_t}$ to the mean $E_t[x]$ of distribution p_t. We can use $\underset{\simeq}{\overline{\widetilde{x}}}\equiv[\underset{\simeq}{x},\overline{\widetilde{x}}]:=\bigcup_t\underset{\sim}{\widetilde{x}_t}$ as a (conservative, but not shortest) robust credible interval, since $p_t(\underset{\simeq}{\overline{\widetilde{x}}})\geq p_t(\underset{\sim}{\widetilde{x}_t})\geq\alpha$ for all t. We can upper bound $\overline{\widetilde{x}}$ (and similarly lower bound $\underset{\simeq}{x}$) by

$$\overline{\widetilde{x}} \;=\; \max_t(E_t[x]+\Delta\widetilde{x}_t) \;\leq\; \max_t E_t[x] + \max_t\Delta\widetilde{x}_t \;=\; \overline{E[x]}+\overline{\Delta\widetilde{x}}. \tag{17}$$

We have already intensively discussed how to compute upper and lower quantities, particularly for the upper mean $\overline{E[x]}$ for $x\in\{\mathcal{F},\mathcal{H},\mathcal{I},...\}$, but the linearization technique introduced in Section 4 is general enough to deal with all in t differentiable quantities, including $\Delta\widetilde{x}_t$. For example for Gaussian p_t with variances σ_t we have $\Delta\widetilde{x}_t=\kappa\sigma_t$ with κ given by $\alpha=\text{erf}(\kappa/\sqrt{2})$, where erf is the error function (e.g. $\kappa=2$ for $\alpha\doteq 95\%$). We only need to estimate $\max_t\sigma_t$.

For non-Gaussian distributions, exact expression for $\Delta\widetilde{x}_t$ are often hard or impossible to obtain and to deal with. Non-Gaussian distributions depending on some sample size n are usually close to Gaussian for large n due to the central limit theorem. One may simply use $\kappa\sigma_t$ in place of $\Delta\widetilde{x}_t$ also in this case, keeping in mind that this could be a non-conservative approximation. More systematically, simple (and for large n good) upper bounds on $\Delta\widetilde{x}_t$ can often be obtained and should preferably be used.

Further, we have seen that the variation of sample depending differentiable functions (like $E_t[x]=E_t[x|\boldsymbol{n}]$) w.r.t. $t\in\Delta$ are of order $\frac{s}{n+s}$. Since in such cases the standard deviation $\sigma_t\sim n^{-1/2}\sim\Delta\widetilde{x}_t$ is itself suppressed, the variation of $\Delta\widetilde{x}_t$ with t is of order $n^{-3/2}$. If we regard this as negligibly small, we may simply fix some $t^*\in\Delta$:

$$\max_t \Delta\widetilde{x}_t = \kappa\sigma_{t^*} + O(n^{-3/2}).$$

Since $\Delta\widetilde{x}_t$ is "nearly" constant, this also shows that we lose at most $O(n^{-3/2})$ precision in the bound (17) (equality holds for $\Delta\widetilde{x}_t$ independent of t).

Robust credible intervals for mutual information. Consider the mutual information defined in (13). The robust credible interval for $\mathcal{I}$ can be estimated as follows.

$$\overline{\widetilde{\mathcal{I}}} \;\leq\; \overline{I} + \overline{\Delta\widetilde{\mathcal{I}}} \;\leq\; I_0 + I_R^{ub} + \overline{\Delta\widetilde{\mathcal{I}}} \;=\; I_0 + I_R^{ub} + \kappa\sqrt{\mathrm{Var}_{t^*}[\mathcal{I}]} + O(n^{-3/2}).$$

Expressions for the variance of $\mathcal{I}$ have been derived in [Hut01]:

$$\mathrm{Var}_t[\mathcal{I}] \;=\; \frac{1}{n+s}\sum_{ij} u_{ij}\left(\log\frac{u_{ij}}{u_{i+}u_{+j}}\right)^2 - \frac{1}{n+s}\left(\sum_{ij} u_{ij}\log\frac{u_{ij}}{u_{i+}u_{+j}}\right)^2 \;+\; O(n^{-2}).$$

Higher order corrections to the variance and higher moments have also been derived, but are irrelevant in light of our other approximations.

9 Conclusions

This is the first work, providing a systematic approach for deriving closed form expressions for interval estimates for the Imprecise Dirichlet Model (IDM). We concentrated on exact and conservative *robust* interval ([lower,upper]) estimates for concave functions $F=\sum_i f_i$ on simplices, like the entropy. For the conservative estimates we used a first-order Taylor series expansion in one over the sample size n and bounded the exact remainder, which widened the intervals by $O(n^{-2})$. This construction may work for other imprecise models too. Here is a dilemma, of course: For large n the approximations are good, whereas for small n the bounds are more interesting, so the approximations will be most useful for intermediate n. More precise expressions for small n would be highly interesting. We have also indicated how to propagate robust estimates from simple functions to composite functions, like the mutual information. We argued that a reduced IDM on product spaces, like Bayesian nets, is more natural and should be preferred in order to improve predictions. Although improvement is formally only $O(n^{-2})$, the difference may be significant in Bayes nets or for very small n. Finally, the basics of how to combine robust with credible intervals have been laid out. Under certain conditions $O(n^{-3/2})$ approximations can be derived, but the presented approximations are not conservative. All in all this work has shown that the IDM has not only interesting theoretical properties, but that explicit (exact/conservative/approximate) expressions for robust (credible) intervals for various quantities can be derived. The computational complexity of the derived bounds on $F=\sum_i f_i$ is very small, typically one or two evaluations of F or related functions, like its derivative. First applications of these (or more precisely, very similar) results, especially the mutual information, to robust inference of trees look promising [ZH05].

Acknowledgements. I want to thank Peter Walley for introducing the IDM to me, Marco Zaffalon for encouraging me to investigate this topic, and Jean-Marc Bernard for his feedback on earlier drafts of the paper.

A Properties of the ψ Function

The digamma function ψ is defined as the logarithmic derivative of the Gamma function. Integral representations for ψ and its derivatives are

$$\psi(z)=\frac{d\ln\Gamma(z)}{dz}=\frac{\Gamma'(z)}{\Gamma(z)}=\int_0^\infty\left[\frac{e^{-t}}{t}-\frac{e^{-zt}}{1-e^{-t}}\right]dt,\quad \psi^{(k)}(z)=(-1)^{k+1}\int_0^\infty\frac{t^k e^{-zt}}{1-e^{-t}}dt.$$

The h function (5) and its first derivative are

$$\begin{aligned}h(u_i) &= (n_i+st_i)[\psi(n+s+1)-\psi(n_i+st_i+1)]/(n+s),\\ h'(u_i) &= \psi(n+s+1)-\psi(n_i+st_i+1)-(n_i+st_i)\psi'(n_i+st_i+1),\end{aligned}$$

For integral s and at argument $u_i^0=\frac{n_i}{n+s}$ and $u_i^0=\frac{n_i+s}{n+s}$ we need ψ and ψ' only at integer values for which the following closed representations exist

$$\psi(n+1)=-\gamma+\sum_{i=1}^n\frac{1}{i},\quad \psi'(n+1)=\frac{\pi^2}{6}-\sum_{i=1}^n\frac{1}{i^2},$$

where $\gamma=0.5772156...$ is Euler's constant. Closed expressions for half-integer values and fast approximations for arbitrary arguments also exist. The following asymptotic expansion can be used if one is interested in $O((\frac{s}{n+s})^2)$ approximations only (and not rigorous bounds):

$$\psi(z+1)=\log z+\frac{1}{2z}-\frac{1}{12z^2}+O(\frac{1}{z^4}),$$

This shows that $h(u_i)$ converges to $-u_i\log u_i$ for $n\to\infty$ (and $u_i\to$const.), i.e. $H(\boldsymbol{u})$ is close to $\mathcal{H}(\boldsymbol{u})$ for large n. See [AS74, Chp.6] for details on the ψ function and its derivatives. From the above expressions one may show $h''<0$.

B Symbols

Symbol	Explanation
δ_{ij}	Kronecker symbol ($\delta_{ij}=1$ for $i=j$ and $\delta_{ij}=0$ for $i\neq j$)
$\imath$,i	Discrete random variable, index/outcome/observation $\in\{1,...,d\}$
d	Dimension of discrete random variable $\imath$
π_i	(Objective/aleatory) probability/chance of i
log	natural logarithm to basis e
x_i,$\boldsymbol{x}$,x_+	Vector $\boldsymbol{x}=(x_1,...,x_d)$, $\quad x_+=x_1+...+x_d$, $\quad \boldsymbol{x}\in\{\boldsymbol{n},\boldsymbol{t},\boldsymbol{u},\boldsymbol{\pi},...\}$
t_i,$\boldsymbol{t}$	Initial bias of i, bias vector
Δ	$=\{\boldsymbol{\pi}:\pi_i\geq 0\,\forall i,\sum_i\pi_i=1\}$ $=\boldsymbol{\pi}$-simplex ($\boldsymbol{\pi}\in\Delta$)
$\Delta_{(e)}$	$=\{\,\boldsymbol{t}\,:t_i\geq 0\,\forall i,\sum_i t_i\overset{(<)}{=}1\}$ $=$ (extended) $\boldsymbol{t}$-simplex ($\boldsymbol{t}\in\Delta_{(e)}$)
$\Delta'_{(e)}$	$=\{\boldsymbol{u}:u_i\geq u_i^0\,\forall i,\sum_i u_i\overset{(<)}{=}1\}=$ (extended) $\boldsymbol{u}$-simplex ($\boldsymbol{u}\in\Delta'_{(e)}$)
s	Magnitude of imprecision ($n'_i=st_i$ is virtual observation #)
$\boldsymbol{D}$	Data/sample $\{i_1,...,i_n\}$
n_i,$\boldsymbol{n}$,n	# of outcomes/observations i, # sample vector, total sample size
$\delta(\cdot)$	Dirac delta distribution $\int f(x)\delta(x)dx=f(0)$
$p(\boldsymbol{\pi}\|\boldsymbol{n})$	$\propto\prod_i\pi_i^{n+st_i-1}\propto$ Dirichlet posterior (second order/belief/subjective/epistemic probability)
$E_{\boldsymbol{t}}[\mathcal{F}]$	Expected value of $\mathcal{F}$ w.r.t. posterior $p(\boldsymbol{\pi}\|\boldsymbol{n})$
w.r.t.	with respect to
i.i.d.	independent and identically distributed
u_i^0	$=\frac{n_i}{n+s}$
u_i	$=\frac{n_i+st_i}{n+s}=E_{\boldsymbol{t}}[\boldsymbol{\pi}]$
u_i^*,t_i^*	Origin for Taylor expansion
σ	$=\frac{s}{n+s}=1-u_+^0=$ Taylor expansion parameter
$O(\sigma^k)$	$f(\boldsymbol{n},\boldsymbol{t},s)=O(\sigma^k)$ $:\Leftrightarrow$ $\exists c\,\forall\boldsymbol{n}\in I\!N_0^d,\boldsymbol{t}\in\Delta,s>0:\|f(\boldsymbol{n},\boldsymbol{t},s)\|\leq c\sigma^k$
$\mathcal{H}(\boldsymbol{\pi})$	$=-\sum_i\pi_i\log\pi_i=$ entropy of $\boldsymbol{\pi}$
$H(\boldsymbol{u})$	$=\sum_i h(u_i)=$ expected entropy (see Eq.(5))
$\mathcal{F}(\boldsymbol{\pi})$	$=$ function of $\boldsymbol{\pi}$ ($\mathcal{F}\in\{\mathcal{H},\mathcal{I},...\}$)
$F(\boldsymbol{u})$	$=$ statistic $E_{\boldsymbol{t}}[\mathcal{F}]$ or general function ($F\in\{H,I,...\}$)
$F\sqsubseteq G$	$:\Leftrightarrow F\leq G$ and $F=G+O(\sigma^2)$, i.e. G is "good" upper bound on F
$\boldsymbol{u}^{\overline{F}}$,$\boldsymbol{t}^{\overline{F}}$	maximize (and $\boldsymbol{u}^{\underline{F}}$,$\boldsymbol{t}^{\underline{F}}$ minimize) $F(\boldsymbol{u})$, $\boldsymbol{t}\in\Delta$, $\boldsymbol{u}\in\Delta'$
$\overline{F}$	$=\max_{\boldsymbol{t}\in\Delta}F(\boldsymbol{u})=F(\boldsymbol{u}^{\overline{F}})=$ upper value of $F(\boldsymbol{u})$, similarly $\underline{F}$
$\underline{\overline{F}}$	$=[\underline{F},\overline{F}]=$ robust/Imprecise interval (estimate) of F
$F_0+F_R(\boldsymbol{u})$	$=F(\boldsymbol{u})$ with $F_0=F(\boldsymbol{u}^0)$ and $F_R(\boldsymbol{u})=O(\sigma)$

$[F_R^{lb},F_R^{ub}]$ $\supseteq [\underline{F}_R,\overline{F}_R] \ni F_R$ (conservative [lower,upper] bound on F_R)
$\widetilde{\underset{\sim}{F}}$ $= [\underset{\sim}{F},\widetilde{F}] =$ credible interval (estimate) of F
u_{ij},u_{i+},u_{+j} joint, row, column marginal
$\mathcal{I}(\boldsymbol{\pi})$ $= \sum_{ij}\pi_{ij}\log\frac{\pi_{ij}}{\pi_{i+}\pi_{+j}} =$ mutual information of $\boldsymbol{\pi}$
$I(\boldsymbol{u})$ $= H(u_{i+}) + H(u_{+j}) - H(u_{ij}) = H_{row} + H_{col} - H_{joint}$
$joint,row,col$ Index for quantities based on joint, row, column marginal distr.

References

[AS74] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions.* Dover publications, 1974.

[Ber01] J.-M. Bernard. Non-parametric inference about an unknown mean using the Imprecise Dirichlet Model. In *Proc. 2nd International Symposium on Imprecise Probabilities and Their Application (ISIPTA-2001)*, pages 40–50, The Netherlands, 2001. Shaker Publishing.

[GCSR95] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis.* Chapman & Hall / CRC, 1995.

[Hal48] J. B. S. Haldane. The precision of observed values of small frequencies. *Biometrika*, 35:297–300, 1948.

[Hut01] M. Hutter. Distribution of mutual information. In *Advances in Neural Information Processing Systems 14*, pages 399–406, Cambridge, MA, 2001. MIT Press.

[Hut03] M. Hutter. Robust estimators under the Imprecise Dirichlet Model. In *Proc. 3rd International Symposium on Imprecise Probabilities and Their Application (ISIPTA-2003)*, volume 18 of *Proceedings in Informatics*, pages 274–289, Canada, 2003. Carleton Scientific.

[Hut07] M. Hutter. On universal prediction and Bayesian confirmation. *Theoretical Computer Science*, 384(1):33–48, 2007.

[Jef46] H. Jeffreys. An invariant form for the prior probability in estimation problems. In *Proc. Royal Society London*, volume Series A 186, pages 453–461, 1946.

[Per47] W. Perks. Some observations on inverse probability. *J. Inst. Actuar.*, 73:285–312, 1947.

[Wal96] P. Walley. Inferences from multinomial data: learning about a bag of marbles. *Journal of the Royal Statistical Society B*, 58(1):3–57, 1996.

[WW95] D. H. Wolpert and D. R. Wolf. Estimating functions of distributions from a finite set of samples. *Physical Review E*, 52(6):6841–6854, 1995.

[Zaf01] M. Zaffalon. Robust discovery of tree-dependency structures. In *Proc. 2nd International Symposium on Imprecise Probabilities and Their Application (ISIPTA-2001)*, pages 394–403, The Netherlands, 2001. Shaker Publishing.

[ZH05] M. Zaffalon and M. Hutter. Robust inference of trees. *Annals of Mathematics and Artificial Intelligence*, 45:215–239, 2005.